LAWRENCE LIVERMORE NATIONAL LABORATORY

COMPUTATION

ANNUAL REPORT 2013

LLNL Associate Director for Computation Dona L. Crawford

Deputy Associate Directors James Brase, Patricia Damkroger, John Grosh, and Michel McCoy

Scientific Editors Punita Sinha and John Westlund

Art Director Amy Henke

Production Editor Deanna Willis

Editorial Assistance Caryn Meissner

Proofreader Pamela MacGregor

Photographers Lee Baker, Laura Schulz, Meg Epperly, and George Kitrinos

Print Production Charlie Arteago, Jr., and Monarch Print Copy and Design Solutions

LLNL-TR-648412

CONTENTS

to TOC PG 2

As Lawrence Livermore National Laboratory engages a broader range of national security issues, Computation provides a correspondingly wider range of scientific expertise, especially in energy, intelligence, and cybersecurity. (Cover design by Amy Henke)

RETURN TO TOC

MESSAGE FROM THE ASSOCIATE DIRECTOR

COMPUTATION provides the intellectual electricity that helps make Lawrence Livermore National Laboratory (LLNL) a thriving and innovative research and development (R&D) institution. Our hardware, software, and discipline experts are integrated into almost every field of research at LLNL. Our in-house programs, software applications, and support initiatives help achieve mission-critical national-security objectives for the U.S. Departments of Energy, Homeland Security, and Defense, as well as other federal and state agencies. We deliver comprehensive support to a wide set of scientific disciplines that include physical sciences (materials modeling, fluid dynamics), biology (genomic sequencing), mathematics (scalable algorithms), and computer sciences (parallel processing, real-time control systems). This year, we have expanded our outreach with high performance computing (HPC) collaborations and made great strides in our capabilities related to mathematical computation, control systems, and computation infrastructure, many of them in concert with partners in industry and academia. The Computation workforce remains highly motivated and passionate about melding computational expertise with cutting-edge science and technology to enable breakthroughs.

The IBM Blue Gene/Q Sequoia system continues to rank No. 1 on the Graph 500 as it has since 2011. Sequoia traversed 15,363 giga edges per second on a scale of 40 graph (i.e., a graph with 2^{40} vertices). Collaborators working with Livermore scientists won the Gordon Bell Prize for simulating clouds of 15,000 collapsing bubbles using 6.4 million threads on Sequoia. The simulation set a record for supercomputing in fluid dynamics, using 13 trillion cells and 14.4 petaFLOPs. The team achieved a 150-fold improvement in resolution and a 20-fold reduction in time to solution compared to previous research. As impressive as these numbers and awards are, Livermore is far more focused on applying supercomputing to the world's most pressing problems in national security, health, energy, climate, and business development. For instance, the potential future applications of the collapsing bubbles work on Sequoia could be as diverse as improving the design of high-pressure fuel injectors and propellers, shattering kidney stones, and destroying tumorous cancer cells.

Computation's collaborations are accelerating innovation by bringing together hardware and software engineers, computational scientists, applied mathematicians, and computer scientists. We push the development of next-generation technology to find solutions to intractable problems in areas such as bioinformatics, energy distribution, and national security. To boost U.S. economic competitiveness, our High Performance Computing Innovation Center (HPCIC) works with industrial partners to develop HPC solutions for reducing prototyping time and expense, optimizing globally distributed manufacturing processes, and extracting strategic insights from oceans of data. This year, the HPCIC signed a Memorandum of Understanding with the Hartree Centre at the Daresbury Laboratory in the United Kingdom (U.K.) for the purpose of helping U.S. and U.K. companies gain technological advantages through advanced computing. The Hartree–HPCIC partnership should yield future technical and business development exchanges between the two parties.

Computation's scientific and technological expertise has advanced to the level where we can now model the human heart, beating in real time, as demonstrated by the joint IBM and Livermore Cardioid code. This state-of-the-art simulation, which replicates the electrophysiology of the heart, provides realistic, time-accurate modeling of heartbeats. These simulations can be used in the development of pharmaceutical drugs, heart monitoring devices, and drug toxicity studies aiding in the fight of the nation's number one killer—heart disease.

The Catalyst system, a Cray CS300 cluster supercomputer, is the latest system to be added to our HPCIC resource catalog. This machine will be shared by Intel, Cray, and LLNL and used to speed technology development in HPC simulation and big data analytics. Catalyst is a departure from classic HPC architectures. Our explorations in persistent-memory technology with increased storage using both volatile and nonvolatile memories will benefit research areas from astrophysics to bioinformatics within the Advanced Simulation and Computing (ASC) Program. Access to very large, low-latency memory on a single node can provide orders of magnitude improvements over calculations that use distributed memory on a computing cluster optimized for some physics codes. We expect to gain insight by combining floating-point-focused capability with data-analysis-producing technologies that could be commercialized in the future.

Computation is leading LLNL's Data Science Initiative to develop expertise in capabilities related to big data. The Laboratory built and honed its expertise in simulating large physical systems over the past 60 years. Now we are coupling that expertise with our growing proficiency in data science. The increasingly pervasive Internet and automated collections of data have resulted in large, complex network systems that include social networks, sensors, seismic monitoring, and many other sources of streaming data that are continuously increasing in volume, velocity, and variety. Intense analytics are required to organize, prioritize, analyze, and utilize the data these systems provide. Existing resources are overwhelmed by the massive amount of data. Therefore, it is becoming increasingly important to develop tools to deal with these problems, especially in emerging mission areas such as cybersecurity, network optimization, atmospheric modeling, bioinformatics, and medicine. LLNL's Data Science Initiative includes a five-year R&D plan that will bring together the complementary capabilities of HPC and large-scale data analytics.

Many new partnerships are forming to better understand patterns in big data. For example, we have partnered with Andrew Ng's machine-learning group at Stanford University. The resulting work will allow for pattern discovery in dynamic video and deep-learning neural network data consisting of 100 billion neural connections. Another partnership is being developed with the University of Illinois at Urbana-Champaign, Georgia Tech, and Rensselaer Polytechnic Institute, called the Extreme Parallel Discrete Event Simulation Consortium. The goal is to produce tools that will be applied to several modeling application areas, including network simulation, power grid simulation, agent-based modeling, supercomputer architecture simulation, and circuit simulation.

Another of the big data arenas is biomedicine, which generates enormous amounts of data at the genomic, translational, proteomic, and epigenomic levels. All of this data must be managed and analyzed to find nuggets of knowledge. The Livermore Metagenomic Analysis Toolkit (LMAT) is a new open source software tool designed to address this challenge with accurate and scalable metagenomic analyses.

Through the FastForward program, which is funded by the Department of Energy's Office of Science and National Nuclear Security Administration, we continue to work with technical experts from seven national laboratories and four companies to accelerate the R&D of critical technologies for extreme-scale computing. Next-generation exascale computing is key to LLNL's mission of ensuring the safety and security of the nation's nuclear stockpile without underground testing. In 2013, we worked with Argonne and Oak Ridge national laboratories to specify system technical requirements for the ASC and Advanced Scientific Computing Research programs that define advanced computers targeting 2017 through 2022. This work, known as CORAL, will result in siting three computing systems—one at each of the three laboratories—that will ultimately provide a 4× improvement on the Scalable Science application benchmarks and a 6× improvement on the Throughput application benchmarks compared to current systems.

LLNL has been involved in co-design-like relationships from the early days of supercomputers. This year, we focused on benchmarking effective compiler support. This work is essential for the portability of our codes, which have many years of effort invested in them. One such application is the Livermore Compiler Analysis Loop Suite (LCALS), a new proxy application culled out of the loops based on multigenerational code bases that are at LLNL.

Computation continues to overcome barriers of extreme parallel scaling with improvements in our solvers, integrating developments in computer science and mathematical algorithms. These systems of equations are used to produce simulations of heterogeneous physical systems. We are developing scalable algorithms of O(N) complexity so we can simulate numbers of atoms that are proportional to the hundreds of thousands of processors in our computers. For example, in 2013 we demonstrated scalability of 101,952 atoms on 23,328 processors using the IBM Blue Gene/Q computer Vulcan. We also pioneered a new way of simulating fluid–structure interactions, and our particle transport codes achieved excellent scaling out to 1,572,864 message-passing interface tasks with 37.5 trillion unknowns and 71 percent parallel efficiency.

To encourage innovative ideas, promote teamwork, and enrich employees' work experience, Computation regularly holds "hackathons"— events that give software developers and computer scientists across the Laboratory 24 hours to brainstorm, develop, and improve software products; hack on code; or explore new applications of potential benefit to research programs. LLNL's programs embrace these events, as many of the projects are precursors to products or codes eventually integrated into regular work projects. Geographic information systems, mobile apps, debugging tools, and workflow management tools have been prototyped during these events. Our latest hackathon in November 2013 had 52 participants working on 27 different projects. These events bring out the enthusiasm and verve that generate great collaboration and innovative ideas. They consistently re-energize employees and improve morale.

2013 has once again proven that Computation provides the foundation and support that are essential for productivity at LLNL. Every scientific discipline and program leverages the tools and expertise provided by our staff. We remain at the forefront of HPC moving towards exascale and have developed the infrastructure and proficiency to make this an achievable goal. Our researchers repeatedly deliver outstanding research, attracting collaborations with the biggest names in academia and industry, while joining forces with other national laboratories to further push the bounds of what is scientifically possible. Computation's mission is echoed by the President of the United States in his assertion that "our future depends on reaffirming America's role as the world's engine of scientific discovery and technological innovation." Driven by this lofty goal, we employ our state-of-the-art computational capabilities to ensure LLNL's position as a global science and technology leader.

Dona Crawford
Associate Director, Computation

AN AWARD-WINNING ORGANIZATION

THE stories in this annual report present a cross section of Computation's accomplishments in research, high performance computing (HPC), software applications, and information technology and security. In addition to the projects highlighted in the report, Computation personnel and projects received prestigious external recognition in 2013. Several notable accomplishments are featured in this section.

A NEW SIMULATION SPEED RECORD

Computer scientists at LLNL and Rensselaer Polytechnic Institute (RPI) set an HPC speed record that opens the way to the scientific exploration of complex, planetary-scale systems. In a paper published in May, the team announced a record-breaking simulation speed of 504 billion events per second on LLNL's Sequoia IBM Blue Gene/Q supercomputer, dwarfing the previous record set in 2009 of 12.2 billion events per second.

In addition to breaking the record for computing speed, the research team set a record for the most highly parallel discrete event simulation, with 7.86 million simultaneous tasks using 1.97 million cores. Discrete event simulations are used to model complex irregular systems with behavior that cannot be described by equations, such as communication networks, traffic flows, economic and ecological models, military combat scenarios, and many others.

Authors of the study are Lawrence Livermore scientists David Jefferson and Peter Barnes, Jr.; and RPI's computer science professor Chris Carothers of the Computational Center for Nanotechnology Innovation and graduate student Justin LaPre.

The records were set using the ROSS (Rensselaer's Optimistic Simulation System) simulation package developed by Carothers and his students, and using the Time Warp synchronization algorithm originally developed by Jefferson.

SEQUOIA TOPS GRAPH 500 LIST OF BIG DATA SUPERCOMPUTERS

Sequoia retained its No. 1 Graph 500 ranking in June and November, showcasing its ability to conduct analytic calculations—or find the proverbial

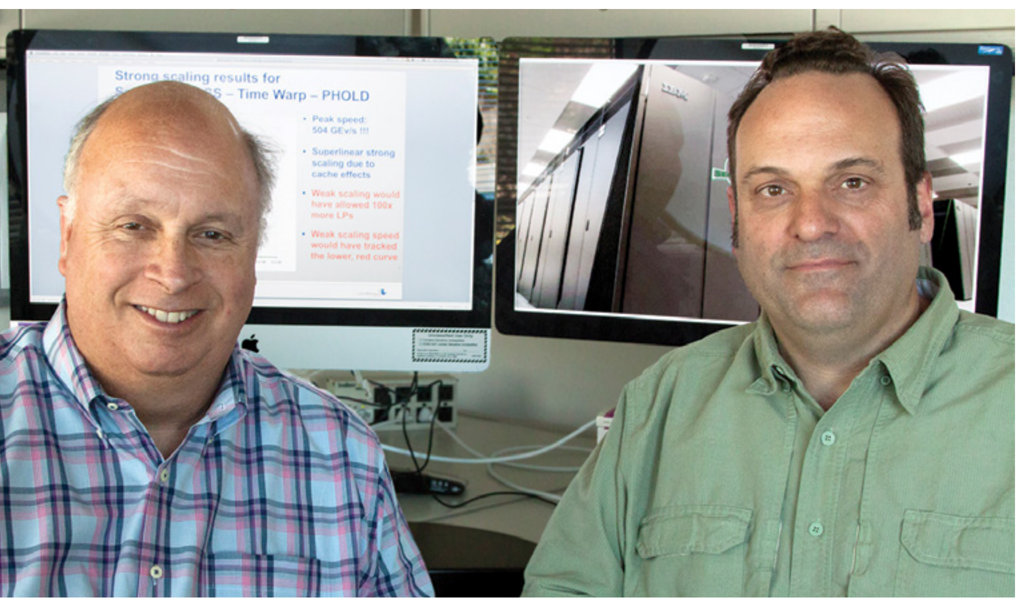

Lawrence Livermore scientists David Jefferson and Peter Barnes, Jr.

The Sequoia supercomputer retained its No. 1 ranking on the Graph 500 list and enabled groundbreaking and award-winning science throughout 2013.

needle in the haystack—by traversing 15,363 giga edges per second on a scale of 40 graph (a graph with 2^40 vertices). The biannual list ranks HPC systems on the basis of processing massive amounts of big data. Graph 500 looks at graph-based data problems, a foundation of most analytics work, and the ability of systems to process and solve complex problems. The Graph 500 was established as a complement to the TOP500 list, and Sequoia has held the top spot since November 2011. This feat reflects Sequoia's ability to push the boundaries of the data-intensive computing that is critical to LLNL's national security missions.

SCIENTISTS "BURST" SUPERCOMPUTING RECORD WITH BUBBLE SIMULATION

Collaborators at ETH Zurich, IBM Research, and the Technical University of Munich working with Lawrence Livermore scientists set a record in supercomputing in fluid dynamics by resolving unique phenomena associated with clouds of collapsing bubbles. The calculation was awarded the prestigious Gordon Bell Prize for peak performance at SC13 in Denver, Colorado in November.

Using 6.4 million threads on Sequoia, the team simulated 15,000 bubbles at 14.4 petaFLOP/s of sustained performance—73 percent of the supercomputer's theoretical peak. When harnessed, the violent bursting of bubbles can shatter kidney stones and destroy tumorous cancer cells. This research could open up new opportunities to collaborate with medical doctors to further explore the applications of bubble cavitation in healthcare.

LLNL computer scientists Adam Bertsch, Blue Gene systems lead, and Scott Futral, group leader for the HPC development environment, assisted the Gordon Bell Prize-winning team. Livermore Computing enabled the achievement of this simulation on Sequoia.

The bubbles simulations are one to two orders of magnitude faster than any previously reported flow simulation. The last major achievement was earlier this year by a team at Stanford University, which broke the one million core barrier, also on Sequoia.

Lawrence Livermore's Adam Bertsch, Dona Crawford, and Scott Futral with the certificate for No. 1 on the Graph 500 in the SC13 Department of Energy booth.

Anna Maria Bailey, HPC facilities manager for LLNL; Tom Tabor, publisher of *HPCWire*; Fred Streitz, HPCIC director; and Becky Springmeyer of the Advanced Simulation and Computing (ASC) Program. LLNL won two *HPCWire* awards, "Best Application of 'Green Computing' in HPC" and "Best HPC Collaboration Between Government and Industry."

RECOGNITION FROM TOP NEWS SERVICE

LLNL was selected by readers of the *HPCWire* news service for a 2013 Reader's Choice Award for "Best Application of 'Green Computing' in HPC."

The Laboratory also received an Editor's Choice Award by editors of *HPCWire* for outreach to industry through the High Performance Computing Innovation Center (HPCIC) and special programs such as the hpc4energy incubator. The award was for "Best HPC Collaboration Between Government and Industry."

Computation Associate Director Dona Crawford was named one of *HPCWire's* "People to Watch" in 2013. This was Crawford's second time receiving this distinction; she was the first woman to be named to the list in 2002, and she is the first woman to have made the list twice. The 12 people selected by *HPCWire* each year are among the most talented and outstanding individuals in the HPC community. Recipients are selected from a pool of potential candidates in academia, government, and industry.

A RECORD-SETTING SUMMER PROGRAM

The Computation Directorate hosted another excellent and productive summer program, welcoming 118 students and faculty—77 graduate students, 34 undergraduates, 1 high school student, and 6 faculty—from 74 universities and 6 countries. This year's group was the most diverse group of students Computation has ever hosted. Students represent the future of computing at LLNL; their success will be our legacy.

Computation's Institute for Scientific Computing Research hosted 118 students and faculty during the 2013 summer program.

1

ADVANCES IN CORE COMPETENCIES ENSURE THE VITALITY OF THE DISCIPLINE

CORE COMPETENCIES

TERRI QUINN
QUINN1@LLNL.GOV

THE Computation Directorate is home to a broad range of computing expertise and capabilities. At our foundation is a world-class staff of experts that help Lawrence Livermore National Laboratory (LLNL) meet its current and future mission objectives. Building on this foundation are our signature core competencies. The Computation Directorate is the steward of three core competencies: computational mathematics, data science, and high performance computing (HPC).

HPC has long been a defining strength of LLNL. Programs have an overarching need to simulate the behavior of complex systems with quantifiable confidence. Accordingly, the Laboratory makes major investments in HPC computers, has considerable expertise in scalable computing environments, and develops new algorithms and state-of-the art simulation codes. In 2013, these investments bore fruit to yield a remarkable computer, Sequoia. This machine is running large modeling and simulation workloads for the nuclear weapons program and can run several days without a failure—impressive for a computer with 96,000 nodes; 1.6 million processing elements; and 1.6 petabytes of memory. Concurrently, we have begun planning our next advanced HPC computer. As with Sequoia, we will work closely with a leading U.S. supercomputing company to develop a machine that will be affordable, sufficiently reliable, and capable of running even more challenging problems than are running on Sequoia.

LLNL's High Performance Computing Innovation Center (HPCIC) has proven to be a productive conduit for collaborating with industry in HPC. The success of the Cardioid code is proof of what record-breaking HPC can do for industry. To serve both industry and LLNL application areas, we recently worked with Cray and Intel to develop an innovative computer called Catalyst, a new breed of computer that combines features commonly found in HPC and large data analytics.

Concurrently, Computation is building on our strength in graph analytics, scientific visualization, machine learning, distributed data management, and persistent memory research to create a coherent data science program that addresses problems in cybersecurity, electric grid optimization, biodefense, and climate change. We are collaborating with Stanford University, Rensselaer Polytechnic Institute, and New York University's Center for Urban Science and Progress to advance data science capabilities.

Computation's internationally renowned expertise in computational mathematics continues to produce new models, algorithms, and software for the predictive simulation of physical phenomena. For example, this year, we demonstrated first-principles molecular dynamics algorithms that scale to more than 100,000 atoms on the IBM Blue Gene/Q Vulcan system for studying matter at atomistic scales by calculating interactions between atoms using quantum mechanical models. This replaces standard N^3 algorithms, which can efficiently scale to only 500 atoms, even on today's largest computers. Without such algorithmic advancements, we cannot fully utilize our HPC machines to deliver groundbreaking scientific insight.

These core capabilities provide the foundation for research and development throughout LLNL and are helping to solve today's problems and prepare for tomorrow's known and anticipated challenges.

CORAL and Sierra Continue Livermore's Computing Legacy

LLNL's HPC strategic objective is to ensure mission success through cutting-edge HPC systems. Sierra will be the next advanced technology system sited at LLNL in the Advanced Simulation and Computing (ASC) Program's system line that has included Blue Pacific, White, Purple, and Sequoia. As the next advanced technology system, Sierra will be expected to address the most demanding computing problems facing the ASC Program. To achieve this goal, the system must provide the largest capability available to ASC applications and incorporate novel technology that foreshadows the future directions of large-scale systems.

Sierra is part of the CORAL procurement, a first-of-its-kind collaboration between Oak Ridge, Argonne, and Lawrence Livermore national laboratories that will culminate in three pre-exascale HPC systems delivered in the 2017 time frame. The CORAL competitive bid process will select two different system architectures—Argonne and Oak Ridge will each receive unique architectures, and LLNL will select one of the two architectures. The Argonne and Oak Ridge systems will help meet the future mission needs of the Advanced Scientific Computing Research (ASCR) Program within the Department of Energy's (DOE's) Office of Science, while Sierra will serve the mission needs of the ASC Program within the National Nuclear Security Administration (NNSA).

Many challenges arise from current computer architecture trends. Thus, the technological advances that are required for a highly usable system will not occur in the CORAL time frame without a deliberate and strategic investment plan. Thus, the CORAL collaboration includes not just the procurement of the systems but also targeted investment in nonrecurring engineering (NRE) research and development contracts. Overall, the CORAL procurement is one of the most complex in DOE's history.

The three laboratories will jointly negotiate the three build contracts and two NRE contracts. The contracts will be focused on the research and development work needed to deliver the three computer systems. The respective laboratories where the systems will be placed will issue the build contracts. LLNL will issue the NRE contracts, which technical representatives from all three laboratories will oversee.

FY13 CORAL activities involved extensive preparations for this complex procurement. The year began with

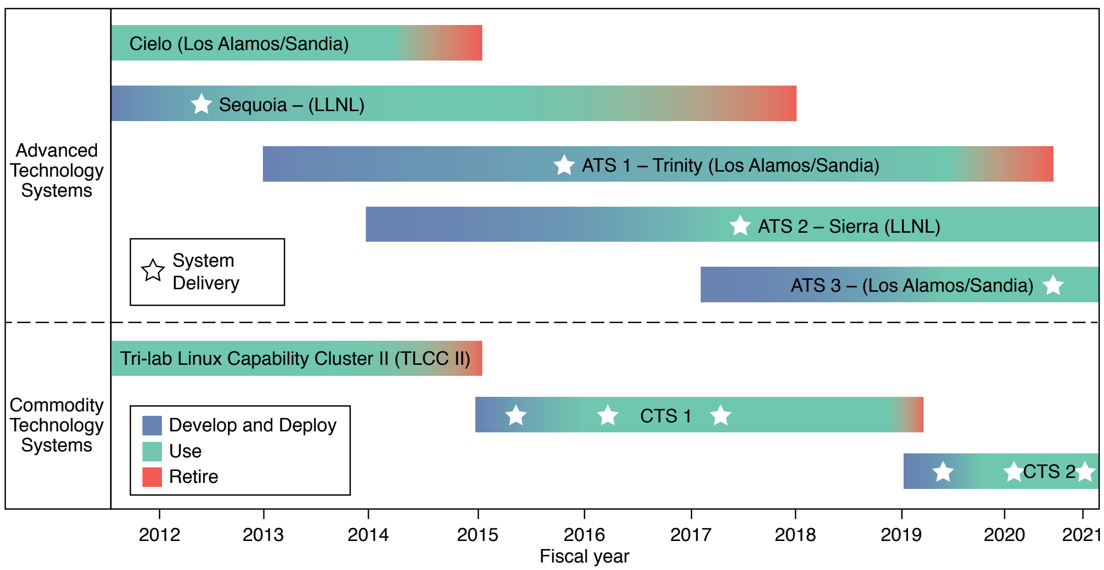

The planning timeline for Advanced Simulation and Computing (ASC) system procurements is divided into two categories: Advanced Technology Systems (ATSs) and Commodity Technology Systems (CTSs). ATSs are intended to address the most demanding computing problems facing the ASC Program, whereas CTSs use mature technology available in the commodity market to address the day-to-day workload of the ASC Program. In general, today's ATSs serve as precursors to tomorrow's CTSs.

BRONIS R. DE SUPINSKI
BRONIS@LLNL.GOV

the release of a Request for Information (RFI) by Oak Ridge on behalf of all three laboratories. The RFI responses, as well as other market survey activities, provided extensive information on the technologies that will be available in the CORAL time frame and their anticipated costs. The laboratories also completed Critical Decision documents that denote the mission needs of the ASC Program and ASCR between 2017 and 2022. A draft CORAL technical requirements document was developed that combines the capability required to meet those mission needs with the information in the RFI responses to provide realistic expectations of the eventual CORAL systems. The draft, which was made available to potential system vendors at a meeting in July 2013, will form the basis of one of the key documents in the CORAL Request for Proposals (RFP), which will be released in FY14.

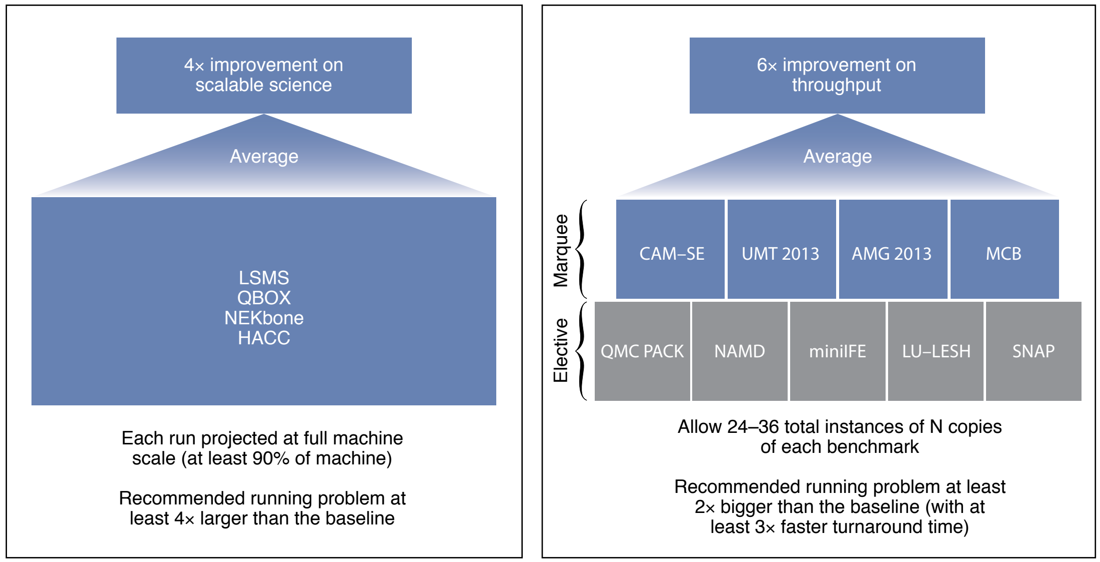

CORAL benchmarks capture key application and workload requirements.

The CORAL benchmarks form the heart of the CORAL procurement strategy. The most important technical requirements for the RFP will be to meet the performance goals captured through these benchmarks. The benchmarks have, therefore, been carefully chosen and developed to represent important aspects of a broad range of applications expected to dominate the science and mission deliverables on the CORAL systems. Reflecting the expected uncertainty quantification and large-scale science workloads of Sierra and the Argonne and Oak Ridge systems, the CORAL benchmarks are divided into two categories, Scalable Science and Throughput.

Scalable Science benchmarks are full applications that are expected to scale to a large fraction of the CORAL system. Typically, these are single-physics applications designed to push the boundaries of human understanding of science in areas such as materials science and combustion. Throughput benchmarks represent specific subsets of applications that are expected to be used as part of the everyday workload of science applications at all three national laboratories. Compared to current systems (Sequoia, Mira, and Titan), CORAL RFP responses are expected to achieve a 4× improvement on the Scalable Science benchmarks and a 6× improvement on the Throughput benchmarks.

The efforts throughout FY13 laid the foundation for the CORAL procurement, which will come to fruition in FY14. The RFP will be released in early FY14, with responses due approximately six weeks after its release. With input from LLNL's ASC Tri-Lab partners, technical experts from the three CORAL laboratories will evaluate the technical quality of each proposal and then select the two proposals that provide the best overall value to DOE. Following negotiations for the three build contracts and the two NRE contracts, we will enter an exciting period of anticipation for these new world-class systems and begin the hard work necessary to ensure the related expectations are met.

The HPC Innovation Center: Build It and They Will Come

JEFF WOLF
WOLF28@LLNL.GOV

Thanks to prescient, mission-aligned investments by NNSA and LLNL, the High Performance Computing Innovation Center (HPCIC) is actively partnering with U.S. companies, providing access to LLNL's unique computational resources and expertise in order to boost economic competitiveness, accelerate advancements in science and technology, and broaden the HPC skills of the nation's workforce. Strong interest from industrial, academic, and HPC ecosystem partners in these collaborations has transformed the HPCIC from a developmental start-up activity to an established laboratory-wide resource with a proven and attractive external engagement model. Cardioid, the breakthrough human heart modeling and simulation code co-developed by LLNL and IBM researchers, demonstrates the high-impact results achievable through HPCIC engagements when industry partners and LLNL's scientists and engineers collaborate and innovate through the medium of HPC.

Offering to help industry solve high-impact problems by employing some of the world's largest supercomputers—and the software and expertise needed to productively apply them—has become an attractive value proposition for LLNL's HPCIC; the challenge is how to broadly and effectively communicate this offer in ways that capture the attention, imagination, and interest of business leaders. Publicity throughout the year highlighting TOP10 systems Sequoia and Vulcan, outreach programs such as the hpc4energy incubator, and a federal call for collaborative projects expanded the HPCIC's marketplace visibility. Conferences, workshops, and an external website presented further opportunities to tell the HPCIC story across the country. Market analyst firm IDC provided independent third-party validation in its technical computing research report on LLNL's expanded focus on industrial partnerships through the HPCIC.

HPC demonstration projects like Cardioid, a high-resolution whole human heart model that can simulate heartbeats with near-real-time performance, help galvanize the vision of industry leaders and expand perspectives on what is possible with extreme levels of computing. In just one year, a team of 12 scientists from LLNL and IBM Research developed this software from scratch, achieving near-cellular-level resolution and performance 1,200 times faster than the previous state of the art. Cardioid can address a wide variety of applications in pharmaceuticals, medical devices, and clinical medicine, including drug screening for cardiotoxicity and new cardiovascular drug development, and ultimately save lives in the fight against heart disease, the nation's number one killer.

The HPCIC also promotes its business model for engagements with industry to thousands of scientists and engineers inside LLNL. In petascale computing, multidisciplinary project teams of 10 or more scientists are common. Coordination of larger HPC projects and the support of industrial proprietary work has required new rules of engagement and more rigorous business development and program management processes. To help meet these challenges, the HPCIC has established and regularly convenes a business advisory council, composed of management representatives from across LLNL. This council provides visibility and a working interface to all areas of the Laboratory to connect industrial interests to LLNL's scientific computing community, approve projects, establish relative priorities, and resolve resource contention.

As connections are made between prospective industry partners and LLNL researchers, conversations quickly get specific as business and technology experts explore, identify, and define opportunities for engagement that offer high return-on-investment potential for companies and strong alignment with LLNL's mission. Many such conversations occurred this year and have led to both turnkey and collaborative, on-demand engagements with industry for the Laboratory through the HPCIC.

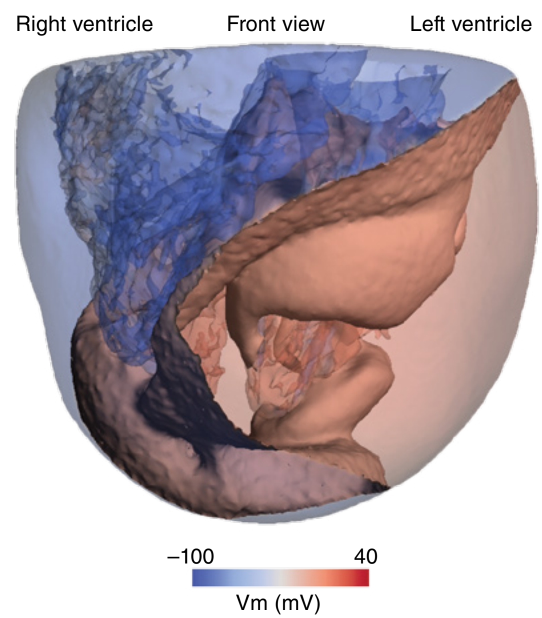

An electrophysiological simulation of reentrant ventricular cardiac arrhythmia shows transmural wave propagation, imaged mid-stroke using Cardioid. Enhanced visualization techniques provide insight into excitations throughout the heart depth using selective cutaways. Visualization by Liam Krauss.

Server Virtualization and System Log Management Improve the Systems Environment

DAVID W. SMITH
SMITH107@LLNL.GOV

Improving the infrastructure systems that support Livermore Computing's (LC's) HPC systems environment is a multiyear effort. Two recent areas of improvement include server virtualization and system log management. Over the past 10 years, LC's infrastructure has matured into a large number of systems that support various services. In a physical server environment, approximately 80 percent of a data center's servers have an average use rate of less than 10 percent. LC's data centers were reaching their limits in terms of floor space, energy, and cooling. Server virtualization addresses these issues and offers many benefits that will improve LC's infrastructure.

We are also centralizing LC's log database environment to improve problem analysis in the HPC clusters, networking infrastructure, and across the computer center. Prior to this effort, logs were analyzed for issues on a system-by-system or network-by-network basis, making correlating events across LC's systems environment an arduous task. This project is expected to enhance our ability to proactively monitor the entire computing environment.

In the past, whenever an LC customer requested a specific service, a server was deployed. This practice led to a proliferation of physical servers and high operational costs given their actual use. LC's ability to respond quickly to customer demands for new or upgraded services was hampered by the service-per-physical-server model.

One of the primary benefits of server virtualization is the decoupling of a specific service or application from a particular piece of hardware (a physical server). Server virtualization removes the physical hardware constraints and allows tailoring of resources. Virtualization technology takes advantage of multicore processors, allowing 10 to 20 virtual servers to run on a single physical server. Consolidating servers at a 20-to-1 ratio has greatly reduced our data center footprint and power consumption.

Prior to 2013, deploying a service involved a lengthy process of server procurement and data center siting (physical racking, network wiring, and electrical work)—a process that was cumbersome and inefficient for the support of our growing number of services and for replacing aging infrastructure. With LC's current deployment of Cisco Unified Computing System (UCS) blade servers to support server virtualization, the physical siting and networking effort has been greatly reduced into a one-time event. The integration of server and networking in the Cisco UCS environment allows for the rapid deployment of virtual servers compared to the service-per-physical-server model.

Also in 2013, LC deployed a centralized log database environment based on Splunk, a tool designed to capture, index, and correlate real-time system events in a searchable repository. Logs and system events are directed from HPC clusters, Lustre file system clusters, network switches, and infrastructure systems and are aggregated into a central database. As queries are developed, correlation of user job failures can be identified and mapped back to system, file system, or network-based events across the computing environment. The Splunk centralized logging system ultimately provides an improved environment for customers.

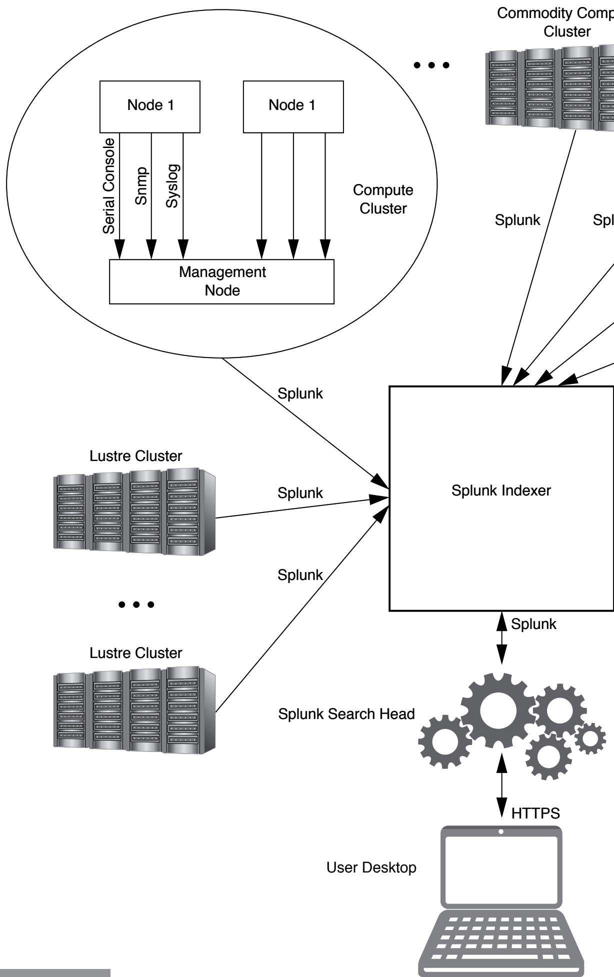

The new Splunk-based centralized logging system indexes and makes searchable data from across Livermore Computing's high performance computing system environment.

Cost-Saving Cooperation Through Archival Storage Quotas

JERRY SHOOPMAN
SHOOPMAN2@LLNL.GOV

One of the unique, core strengths of LC's data storage infrastructure is the world-class digital archives provided by the High Performance Storage System (HPSS). HPSS archives provide high-speed, long-term retention of data products representing more than four decades of scientific investigation and simulation. To control the cost of long-term storage for our programmatic customers, LC developed an archival quota (aquota) system in 2011. This cooperative approach became especially important in 2013 with the classified deployment of Sequoia, a supercomputer capable of generating massive amounts of data and with the potential to overwhelm LC storage capabilities. Since its initial deployment, aquota has helped LC avoid more than $5M of archive tape media costs.

Several years ago, during the early planning stages for Sequoia, LC staff recognized that archive growth was surpassing historical trends and more than doubling every year. This growth was directly related to the memory size of the platforms. With Sequoia's anticipated 1.6PB of memory, it was clear that demand for storage would be unsustainable. Capacity projections showed LC would need more data center space than was available to store and manage the deluge of data from Sequoia.

In targeted cooperation with several of the largest users of the HPSS archives, LC's Data Storage Group developed and delivered an advisory aquota system in 2011. The aquota system provides oversight and management capabilities for monitoring and controlling archive growth by the largest users. It also provides programmatic management control to allocate additional space as needed. The system consists of client/server software, a MySQL database with access to daily HPSS accounting records, and a graphical plugin on the MyLC webpage. The plugin was newly enhanced with additional graphical detail in 2013. The aquota concept is simple: provide users and programs with a target amount of new data they can store at the beginning of each fiscal year, and use programmatic budgets to determine the amount of growth allowed for the LC user base each year. The system is only concerned with new data because the size of old data becomes inconsequential.

In 2013, the aquota system proved its value. If LC archives had continued to grow at the rate of 2.16× platform memory as they were during early Sequoia planning phases, the archives would today be more than 120PB. Instead, they are 57PB. Without aquota, LC would have needed to purchase at least 17,000 additional tape cartridges at a cost of more than $5M. In addition, the growth constraint provided by aquota in cooperation with archive users has positioned LC to consolidate the archive tape library footprint. In 2013, LC took several large pieces of storage hardware off maintenance plans, achieving an expected annual savings of more than $58K beginning in FY14.

Ultimately, the aquota system has enabled LC to provide a world-class archive at an affordable cost point without large-scale budget increases and without sacrificing the long-recognized reliability that digital tape is known for, all while maintaining extreme energy-efficiency standards. With HPSS archives that are now smaller, more cost-effective, and growing at a controlled pace, LC is better positioned for the challenges and demands of exascale systems.

Livermore Computing (LC) robotics expert Geoff Cleary calibrates a robotic handbot in an SL8500 tape library with 10,000 tape cartridge slots. An LC-developed archival quota system has resulted in a smaller archive footprint as well as a significant cost-savings.

Catalyst Advances Big Data, Simulation, and Technology Partnerships

MATTHEW LEININGER
LEININGER4@LLNL.GOV

LLNL has several mission areas, including stockpile stewardship, bioinformatics, nonproliferation, counterterrorism, as well as broader national security programs. Each of these mission areas requires data science and analytics capabilities in addition to, or in conjunction with, LLNL's HPC simulation capabilities. In collaboration with two U.S. industry partners, Intel and Cray, LLNL has deployed a one-of-a-kind system, Catalyst, to explore the system software and application algorithms for addressing future data analytics and modeling issues. Through LLNL's HPCIC, other U.S. companies can also use Catalyst to investigate and develop solutions for industrial applications.

For more than a decade, LC has successfully pushed the state of the art in HPC cluster technologies by transforming Linux clusters from small departmental systems into large-scale workhorse platforms that are capable of running our most challenging scientific and engineering simulations. Today, these simulations can generate terabytes to petabytes of data, thereby creating a "big data" analysis problem. Big data can also arise in data collection from experiments, sensor or social networks, biological system studies, and many other sources. To explore the future needs of HPC and big data, LLNL, Intel, and Cray created a unique system called Catalyst.

The 150-teraFLOP/s Catalyst cluster has 324 nodes, 7,776 cores, and employs the latest generation Intel 12-core Xeon E5-2695v2 processors. Catalyst runs the NNSA-funded Tri-Laboratory Open Source Software that provides a common user environment across clusters at Livermore, Sandia, and Los Alamos national laboratories. Catalyst boasts 128GB of dynamic random access memory per node, 800GB of nonvolatile memory (NVRAM) per compute node, 3.2TB of NVRAM per Lustre (a type of parallel distributed file system) router node, and improved cluster networking with dual rail Quad Data Rate (QDR-80) Intel TrueScale.

The Catalyst architecture is expected to yield insight into the technologies the ASC Program will require over the next 5 to 10 years to meet high performance simulation and big data computing mission needs. The increased storage capacity of the system (in both volatile and nonvolatile memory) represents the major departure from classic simulation-based computing architectures common at DOE laboratories and opens new opportunities for exploring the potential of combining floating-point-focused capability with data analysis, in one environment. Consequently, the insight provided by Catalyst could become a basis for future commodity technology procurements.

LLNL's partners, Intel and Cray, will take advantage of this one-of-a-kind system to explore different approaches and algorithms for addressing various data analytics and modeling issues both inside and more tangential to the DOE mission space. At LLNL, the HPCIC will offer access to Catalyst and the expected big data innovations it enables as new options for the center's ongoing collaborations with American companies and research institutions. Catalyst will extend the range of possibilities for the processing, analysis and management of the ever-larger and more complex data sets that many areas of business and science now confront. These essential collaborations accelerate innovation by bringing together hardware and software engineers, computational scientists, applied mathematicians, and computer scientists to develop next-generation solutions for pressing problems in such diverse areas as bioinformatics, energy distribution, and national security.

Catalyst is a unique high performance computing (HPC) cluster that will serve research scientists and provide a proving ground for new HPC and big data technologies and architectures.

Data Science Initiative Unites Analytics and Simulation for National Security Applications

Jim Brase
BRASE1@LLNL.GOV

National security priorities are increasingly focused on the resilience and security of complex networked systems such as the Internet, the national power grid, and the global shipping network. Applications and programs in counterterrorism, nuclear weapons nonproliferation, critical infrastructure protection, and cybersecurity are central to the national security missions of the U.S. government.

Today, our capability to discover patterns in data from these large-scale systems, build models of their behaviors, and predict their response to natural disturbances and man-made attacks is limited. Developing analytic and simulation capabilities for these complex socio-technical systems requires new computational approaches that bring together data analytics and simulation in new and fundamental ways.

Over the past 50 years, the national security and scientific communities have established a robust capability to simulate complicated physical systems. These simulations are based on solving sets of partial differential equations that describe the system. Although these applications remain important, emerging priorities involve complex, interconnected socio-technical networks with both continuous and discrete components that are not easily described by sets of equations. Instead, significant components of these models must be based on patterns learned from data.

"Big data" analytics have emerged as a driver for leading-edge computing technology, building on the simultaneous transition to high performance multicore servers. Until now, big data applications have run almost exclusively on commodity clusters with relatively low-performance interconnects and programming environments, such as MapReduce and Hadoop, aimed at data parallel applications. As increasingly complex analytic questions are introduced, computing requirements will begin to overlap the requirements for scientific HPC. The leading edge of analytics—for example, large-scale deep-learning neural networks—has already exceeded the performance offered by commodity clusters. Thus, the trend toward HPC technologies—particularly fast interconnects and accelerators—to overcome communication and reliability limitations will grow and drive the integration of analytic and simulation computing technologies.

LLNL is investing in research projects and developing partnerships to advance this transition. For example, we are demonstrating the potential for data analytics to drive simulations by integrating large-scale computer network mapping with simulations of network activity.

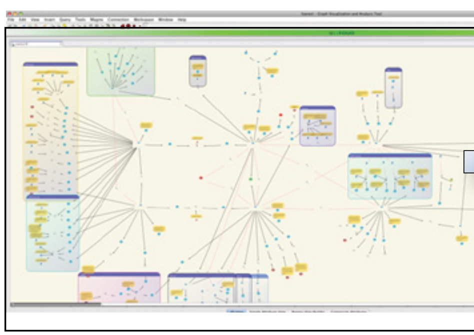

Continuous network mapping: analytics for structure and behaviors

Large-scale network simulation to predict response of the real network

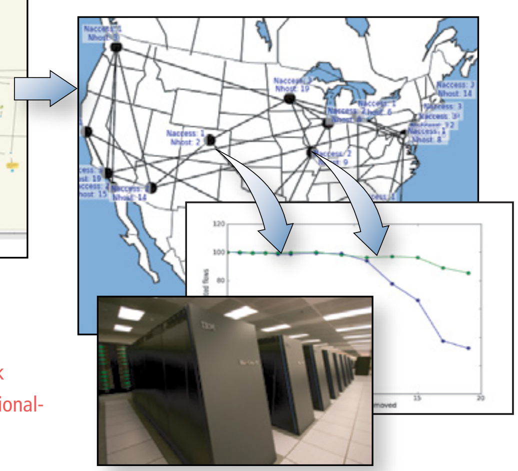

Large-scale analytics are being integrated with network simulation to model the impact of cyber attacks on national-scale networks.

The network structure and activity is measured by mapping systems and then used directly by the simulation, allowing users to explore changes to the network structure and traffic patterns.

In partnership with Stanford University, we are exploring the scaling of deep-learning neural networks to large HPC systems. This research will extend well beyond the current state of the art, creating networks of 100 billion neural connections to enable pattern discovery in dynamic video and network data streams.

LLNL is also partnering with Rensselaer Polytechnic Institute, the University of Illinois, and Georgia Tech to develop new discrete-event simulation methods that, running on LLNL's largest computer systems, can simulate the interactions of billions of agents representing large, complex networks. These innovative tools will allow us to simulate the behaviors of interconnected information and critical infrastructure networks at global scales for the first time.

Integrating data analysis and simulation will enable new classes of science. In particular, urban science and informatics, as embodied at the Center for Urban Science and Progress at New York University, will provide an ideal test bed for these computing approaches and, ultimately, help grow computational social science.

The emergence of private-sector drivers for high performance analytics will dramatically change the HPC landscape. Targeted investments that unite analytics and simulation will support the emerging national security applications of predictive analytics for complex systems.

New Algorithm Enables First-Principles Molecular Dynamics for 100,000 Atoms and Beyond

JEAN-LUC FATTEBERT
FATTEBERT1@LLNL.GOV

First-principles molecular dynamics is a general and fundamental predictive tool, based on quantum mechanics modeling, to calculate interactions between atoms. Unlike classical physics models, where the number of variables (such as temperature and pressure) does not grow with the system size, the number of fields in quantum mechanics models—electronic wave functions—is proportional to the system size. From a computational point of view, quantum mechanics models allow for $O(N^2)$ degrees of freedom for N atoms and $O(N^3)$ floating-point operations for standard solvers. By using such traditional algorithms, one can only produce simulations of a few hundred atoms. Much larger simulations are required to model heterogeneous systems or realistic materials with interfaces, irregularities, and defects. Algorithms with reduced complexity and better parallel scaling are essential to enable simulations with tens of thousands of atoms or more and fully realize the potential of modern high performance computers.

To fully utilize today and tomorrow's computer systems, we are developing a truly scalable algorithm with O(N) complexity and short-range communications, so that one can simulate a number of atoms directly proportional to the number of processors available, for hundreds of thousands of atoms using hundreds of thousands of processors. Our approach relies on the existence of a solution to the density functional theory equations in the form of a set of localized functions, which is possible for insulators with a finite band gap. Numerically, we represent these functions on a uniform finite difference mesh. By confining each function to a specific finite spherical region, one can obtain a very good approximation to this solution with O(N) degrees of freedom using previously developed techniques. Coupling between these localized functions is reduced to the computation of selected elements of the inverse of the Gram matrix, consisting of the dot products between these functions (a particular case of single-particle density matrix).

We use the exponential decay properties of the off-diagonal elements of the inverse of the Gram matrix to compute an accurate approximation of the needed matrix elements in O(N) operations. Our approach is based on a parallel approximate inverse strategy, distributing the computation of the selected elements of the inverse and using only the nearest matrix elements in the Gram matrix to compute each column of the inverse. An efficient parallel communication strategy is used to gather matrix elements from adjacent processors and form local block matrices. Local columns of the inverse of the block matrices are computed using an ILU0 preconditioned generalized minimal residual (GMRES) solver.

In FY13, we demonstrated scalability up to 101,952 atoms on 23,328 processors using the IBM Blue Gene/Q computer Vulcan. We applied our algorithm to a polymer model with a unit cell containing 472 atoms. We then replicated this system by a factor of 2, 3, 4, 5, and 6 in each three-dimensional direction to generate larger simulations and study parallel scaling. We scaled the number of atoms and the number of processors at the same rate, that is, with a constant number of message-passing interface tasks per atom, and achieved excellent weak scaling. Wall clock times for one molecular dynamics step are on the order of one minute. For this particular application, each processor must communicate within a neighborhood of 9×9×5 processors (subdomains are not cubic) to gather and invert matrix blocks with approximate sizes of 2,400.

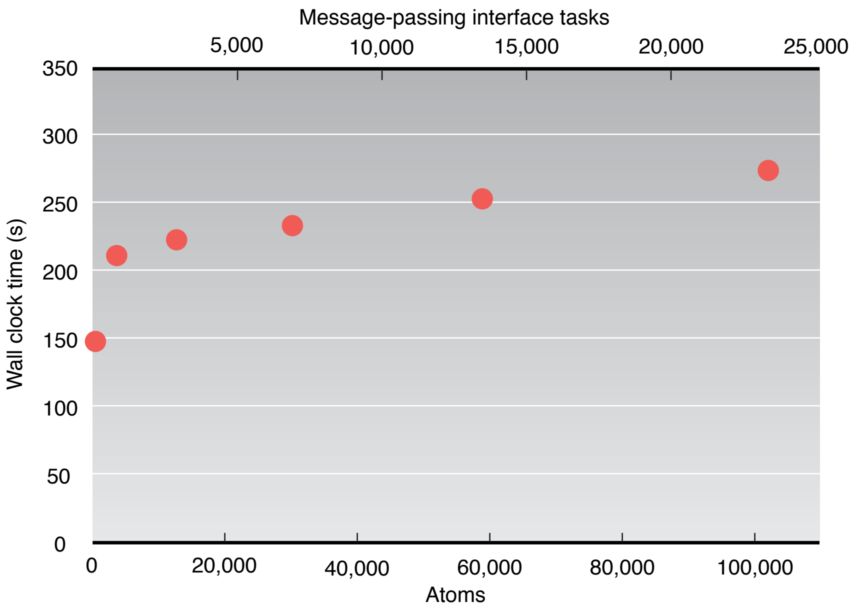

Parallel weak scaling results show the wall clock time for one molecular dynamics step on the IBM Blue Gene/Q architecture and on an Intel Xeon EP X5660 Linux cluster with high-speed interconnect (InfiniBand quad data rate QLogic).

Mathematical Analysis Resolves Important Open Problems in Fluid–Structure Interaction

JEFFREY W. BANKS
BANKS20@LLNL.GOV

WILLIAM D. HENSHAW
HENSHAW1@LLNL.GOV

From designs of nanodevices to nonproliferation policy, technical and policy decisions are becoming increasingly reliant on the simulation of complex, real-world multiphysics systems involving multiple interacting domains. As the power of computers has grown, deficiencies associated with traditional low-order accurate mechanisms for interdomain coupling have become evermore apparent.

Computation researchers are tackling this problem in the context of fluid–structure interaction (FSI) and are developing simulation techniques based on a novel and rigorous mathematical approach. In the past year, innovative techniques for both incompressible and compressible flow have been developed, enabling accurate and efficient simulation of previously inaccessible systems.

The simulation of interacting fluids and solids has proven challenging for various reasons, and a survey of the literature indicates that the presence of so-called added-mass effects has been one of the most significant hurdles. Difficulties associated with the presence of strong added-mass effects have their origin in the fact that the reaction of an immersed body to an applied force depends not only on the mass of the body but also on the mass of the fluid displaced by the body through its motion. For example, a feather or leaf falling in air experiences significant added mass and achieves nothing near the descent rate expected through the application of the force of gravity alone.

FSI algorithms often experience strong added-mass effects as numerical instabilities or nonconvergent iterations. Traditional interface approaches do not properly account for the coupling between fluids and light solids and can, therefore, experience a situation where the overreaction of a light solid to an applied fluid force leads in turn to an even larger reaction from the fluid and so on. LLNL has pioneered a new approach to FSI where a detailed analysis of the evolution of the fully coupled problem exposes stable discrete interface treatments and thereby avoids challenges that have plagued numerical simulation of FSI problems for decades.

The FSI regime coupling high-speed compressible flow and rigid bodies, important in the prediction of blast wave effects on structures, is particularly useful in illustrating some of the difficulties associated with strong added-mass effects. Consider an example consisting of a rigid body obeying Newton's law of motion, embedded in a compressible fluid. A massless body would apparently experience infinite acceleration for any applied fluid force. However, using the new approach, it becomes clear that the problem is dominated by the fluid added mass. This insight was used to develop a partitioned FSI solver that is provably stable even for the case of zero-mass bodies. In addition, the analytic forms for the added-mass tensors for compressible flow are naturally revealed, a question that had remained unanswered since the early 20th century during the times of Love and Taylor.

2013 has also seen preliminary developments for the case of incompressible fluids that represent a significant and nontrivial leap forward, promising orders-of-magnitude speedup over traditional approaches. For incompressible flows and elastic solids and shells, new provably stable and fully second-order techniques were developed. Furthermore, a detailed analysis showed that the extreme difficulty experienced by traditional coupling is a reflection of the fact that the simple traditional approach may appear useable at low resolution but completely fails at higher fidelity. These exciting new results push the boundaries of simulation capability and keep LLNL at the forefront of FSI research.

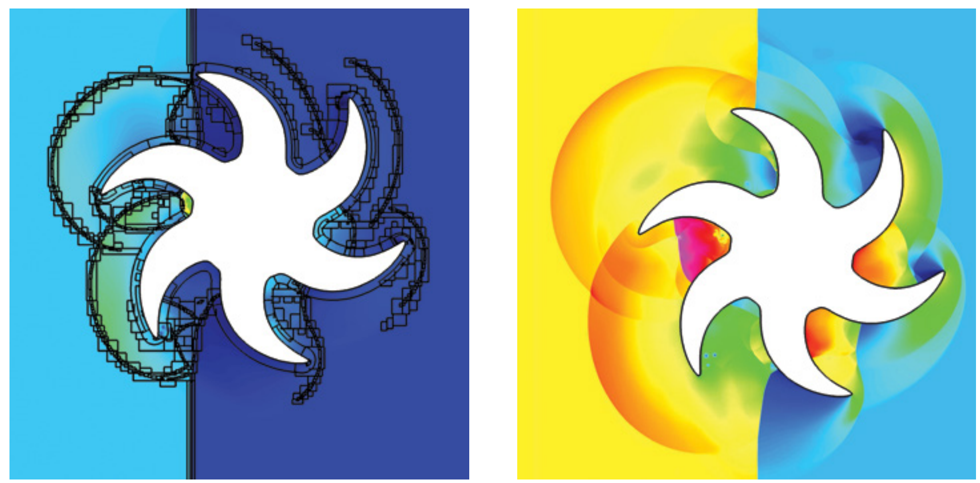

These images from a simulation of shock interaction with a massless rigid starfish use adaptive mesh refinement and deforming composite grids. This simulation would not be possible with existing partitioned approaches.

2 APPLICATIONS AND RESEARCH

COMPUTER APPLICATIONS AND RESEARCH DELIVERS SOLUTIONS TO THE NATION

PEG FOLTA
FOLTA2@LLNL.GOV

LAWRENCE Livermore National Laboratory (LLNL) hosts a vast array of programs aimed at addressing broad national security challenges, from managing the nuclear stockpile to defining new energy sources, developing methods for effective vaccines, understanding climate change, and securing the nation's computational infrastructure. To meet these needs, Computing Applications and Research (CAR) partners with programs to deliver mission-relevant scientific computing and information technology solutions through high-quality software. We seek out and retain a workforce with a unique combination of technical expertise and a strong sense of dedication to the nation.

Issues of scaling remain a focus area, resulting in part from the exponential growth in available scientific data, the multidimensional solution space for such information, and the advances in supercomputing available to support these new solutions. As an example, LLNL's computational experts are working hard to scale a key sweep algorithm used in neutron transport codes to petascale and exascale systems in support of Department of Energy (DOE) and National Nuclear Security Administration (NNSA) applications. Recent studies on LLNL's Sequoia supercomputer have demonstrated the improved code achieves excellent scaling and high efficiency.

Other highlighted work is directed at bridging the gap between higher supercomputer peak performance and applications' ability to utilize that performance. The Livermore Compiler Analysis Loop Suite is a proxy application that enables compiler optimizations through collaborations with vendors. LCALS indicates specific areas that hinder compiler performance, paving the way for new optimization benchmarks that may help solve long-standing compiler issues. These improvements not only benefit our DOE missions, but advance private sector missions as well. Additionally, our efforts to improve ensemble-of-models-based uncertainty quantification (UQ) is making petascale UQ studies possible to the benefit of DOE research efforts including inertial confinement fusion, climate change, weapons, and energy.

LLNL's computational expertise is also playing a role in bioinformatics, enterprise-wide physical security, and large-scale computer controls. Using open source software, we have developed a toolkit that enables more accurate, scalable metagenomic analysis and sequencing for biosecurity applications. Earlier work in this area was seeded by the Laboratory's institutional research investment program, leading to advances in metagenomic analysis that have accelerated organism classification. Computation personnel embedded in our Security Organization provided support to the Y-12 National Security Complex after a physical security breach at the facility in July 2012. This collaboration led to enhancements in the site's overall security posture. We have also extended traditional network monitoring tools for the National Ignition Facility's (NIF's) Integrated Computer Control System, allowing us to locate relevant events within millisecond-precision network packets.

Computation retains a vibrant set of employees with expertise in mathematics and computer science to solve problems of national interest. Whether they choose to focus their career in a specific area or work across multiple disciplines and programs throughout their careers, the extraordinary capabilities of our staff are what make CAR and the Laboratory's computational endeavors so successful.

Neutral Particle Transport Code and Sweep Algorithm Scale to Petascale Architectures

Many DOE, NNSA, and LLNL scientific and engineering applications require the ability to model the behavior of neutral particles, such as neutrons and photons, as they pass through and interact with a material. Examples include reactor and shielding design, medical diagnostics development, illicit nuclear material detection simulations, and high-energy-density physics research. Ensuring that a key algorithm used in several particle transport codes scales to more powerful supercomputers is central to the missions of the Laboratory and DOE. LLNL researchers are developing a neutral particle transport code and related sweep algorithms for use with peta- and exascale computing systems. The LLNL team performed a scaling study of the code on the IBM Blue Gene/Q supercomputer Sequoia and achieved excellent scaling to 1,572,864 message-passing interface (MPI) tasks with more than 37.5 trillion unknowns and 71 percent parallel efficiency.

The number of processors within a typical supercomputer is expanding rapidly. Application performance gains on supercomputers, now and in the foreseeable future, largely stem from code modifications that enable a system to complete more computational tasks in parallel, thus making more efficient use of its processors. A crucial step for LLNL computational experts is ensuring that the codes have a scalable solution for the deterministic neutron transport equations modeling neutron and photon behavior. This system of equations is typically a form of the Boltzmann equation, and has a six-dimensional (6D) phase space consisting of energy (1D), direction (2D), and physical space (3D). Deterministic transport solution methods discretize this phase space, resulting in systems having a finite number of energies, directions, and spatial zones, but an extremely large number of solution unknowns.

When solving these equations, LLNL transport codes currently rely on an efficient iterative solution method based on sweeping. A "sweep" is the solution of a lower triangular linear system in which solution data are calculated sequentially, "sweeping" from one side of the

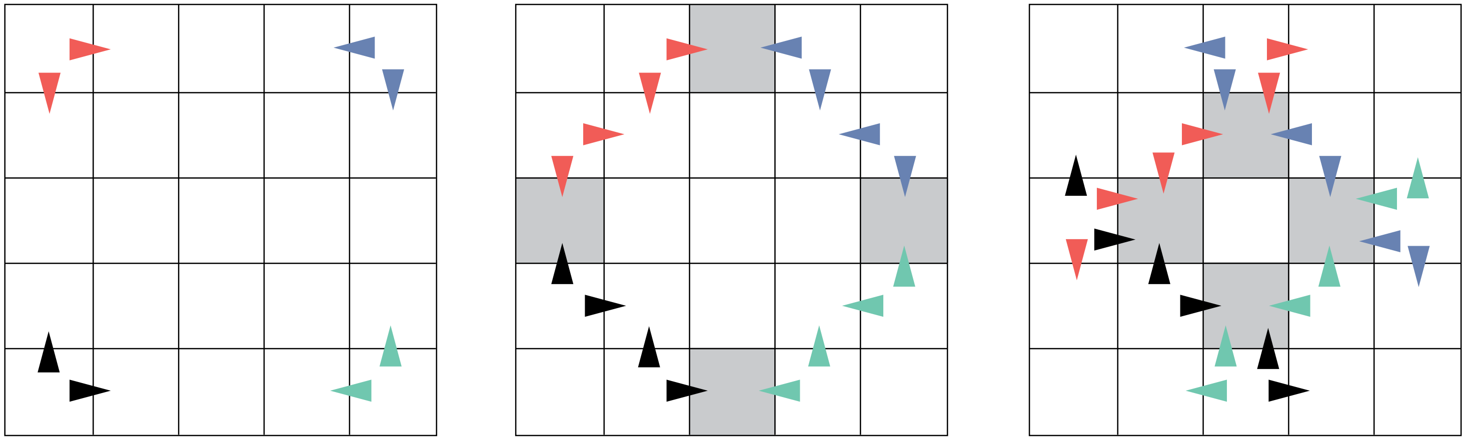

A "sweep" moves data across a spatial mesh for each discrete direction (colored arrows represent the different directions). In the three stages shown here, simultaneous sweeps performed in parallel proceed across subdomains (individual boxes) and pass data from one subdomain to its neighbors. Gray boxes indicate subdomains where data arrive from multiple directions at the same time, forcing the algorithm to choose the order in which the directions are solved. The optimal choice orders the directions from largest to smallest according to how many stages a direction needs to complete its sweep across the domains.

PETER N. BROWN
PNBROWN@LLNL.GOV

spatial mesh to the opposite side. With this technique, data move across the spatial mesh for each discrete direction. One significant scaling advantage of sweep algorithms is that the number of iterations needed to solve the system is determined by the physics of the actual problem and therefore does not increase as the number of unknowns in the problem grows. However, since sweeps have some inherently sequential steps, many researchers have doubted whether they would successfully scale to machines with hundreds of thousands or even millions of processors.

Each iterative step of a solution method requires several sweeps. The total number of sweeps is determined by the product of the number of discrete directions and energies in the problem being solved. As a result, sweep-based algorithms would be intractable for large-scale computing except that each sweep is independent. Parallel transport sweep algorithms take advantage of this autonomy by performing several sweeps simultaneously, decomposing the spatial grid into subdomains and passing information from one subdomain to its neighbors in a data-driven manner. In earlier work, LLNL researchers defined the key components of sweep algorithms that make them efficient and developed parallel performance models of these optimized sweep algorithms. One key aspect of this work involved handling conflicts that arise when data from different discrete directions arrive simultaneously at a subdomain. Several approaches were studied, one of which had already been implemented in the code. An optimized solution was also executed. Contrary to current research, these performance models predicted excellent scaling of the algorithms to petascale architectures.

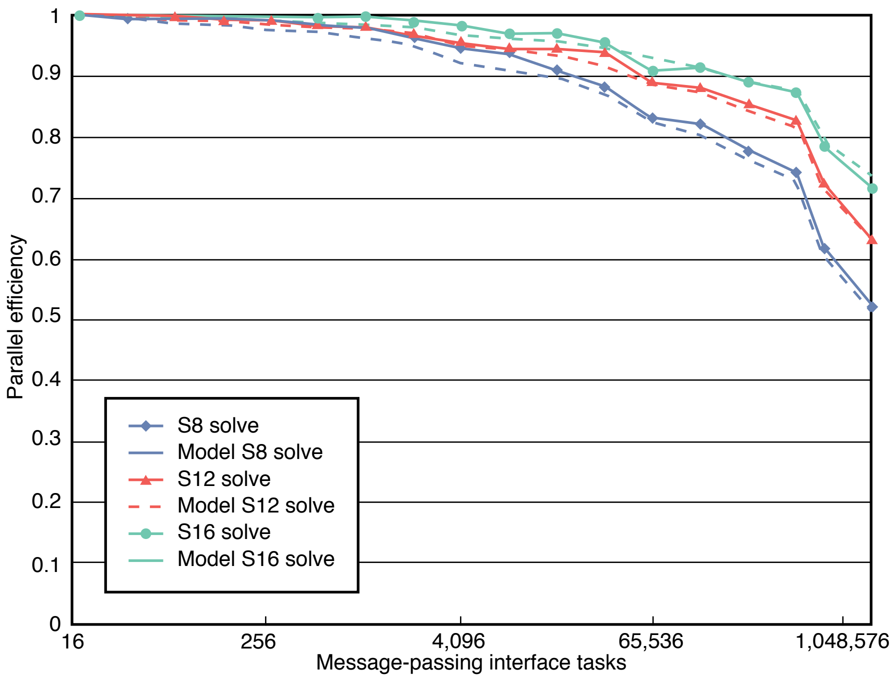

Results of a scaling study on Sequoia compare total solve times for the S8, S12, and S16 discrete directions. Solid lines are code results and dashed lines are model predictions.

Once Sequoia became available for testing, the LLNL team performed a weak scaling study of the code using an earlier version of the optimized approach for resolving data conflicts. The test used the Jezebel criticality benchmark verification problem in which a 6.384-centimeter radius ball of plutonium with a density of 15.61 grams per cubic centimeter is surrounded by air that fills a box measuring 10 centimeters on a side. The problem was run using 48 discrete energies with 80 (S8), 168 (S12), and 288 (S16) discrete directions. The weak scaling study started with one node of Sequoia, using 16 MPI tasks, 1 task per processer, and an initial 48×24×24 3D spatial mesh. The study also used 16-way parallelism in energy, resulting in the best overall scaling. The spatial mesh was scaled to the number of tasks, therefore keeping the problem size per node fixed. The code achieved excellent scaling out to 1,572,864 MPI tasks with more than 37.5 trillion unknowns and 71 percent parallel efficiency (the ratio of the solution time using 16 MPI tasks over that for more MPI tasks). The study proved that sweep algorithms scale much further than previously theorized.

Deep Packet Visualizations Explain Control System Events

MIKHAIL FEDOROV
FEDOROV2@LLNL.GOV

The Integrated Computer Control System for the NIF includes 66,000 devices running on more than 1,900 servers and embedded computers. The devices, which are operated from a central control room, interact over a Local Area Network (LAN) through a middleware framework based on Common Object Request Broker Architecture (CORBA). Having the full details of control system interactions on LAN opens up attractive opportunities for monitoring and debugging the system. While tools are readily available to capture the network traffic, application level analysis is difficult due to the vast size of the raw data. To understand events in the NIF control system that deviate from normal operations, we developed a flexible toolkit that uses Wireshark for deep packet inspection and the R programming language for data analysis and visualizations. The toolkit helped resolve several issues that escaped traditional debugging techniques.

At the network infrastructure level, the toolkit uses the port-mirroring feature of our routers to redirect interesting traffic to an appliance (for example, OPNET ACE Live) or to open source software. After the packets are captured, they are analyzed offline with the Wireshark deep packet inspection tool, which is capable of recovering high-level protocol information from raw capture files. In the case of CORBA, the object method names, object references, and parameter values are recoverable. Given that the packet captures are time-stamped with millisecond precision, an accurate application-level trace of a network event can be reconstructed after the fact.

Finding relevant events in the captured data often presents the biggest challenge. While both the proprietary (OPNET ACE Analyst) and open source (Wireshark) tools come with filters and graphs designed for typical network administration tasks, these tools can be very slow and require multiple manual steps for identifying events in the application domain.

To address the need for problem-specific event finding and visualization tools, we used R language in the toolkit for statistical computing and graphics. The R language is open source software and is designed to handle large data sets with excellent and highly customizable visualization capabilities. The toolkit relies on the tshark utility from the Wireshark package and the ggplot2 package from CRAN (Comprehensive R Archive Network).

In its first use, the toolkit was applied to investigate performance degradations in a system with a distributed controller. The controller worked correctly, but with intervals of extremely slow performance, where the network response times were tens to hundreds of times slower than usual. Increased transmission control protocol (TCP) retransmissions were also seen during the slowdowns. While a normal transaction time is about 1 millisecond, the slowdowns manifested over the scale of hours and days. Hundreds of thousands of operations were involved, most of them showing normal behaviors even during the slowdowns. With the toolkit, the call durations were plotted, partitioned by the call type, and the TCP retransmit overlay was added. Once visualized, the slowdown patterns emerged. With the help of additional informative and intuitive R/ggplot2 visualizations, the root cause was identified. A software fix was developed, and the toolkit confirmed resolution of the problem.

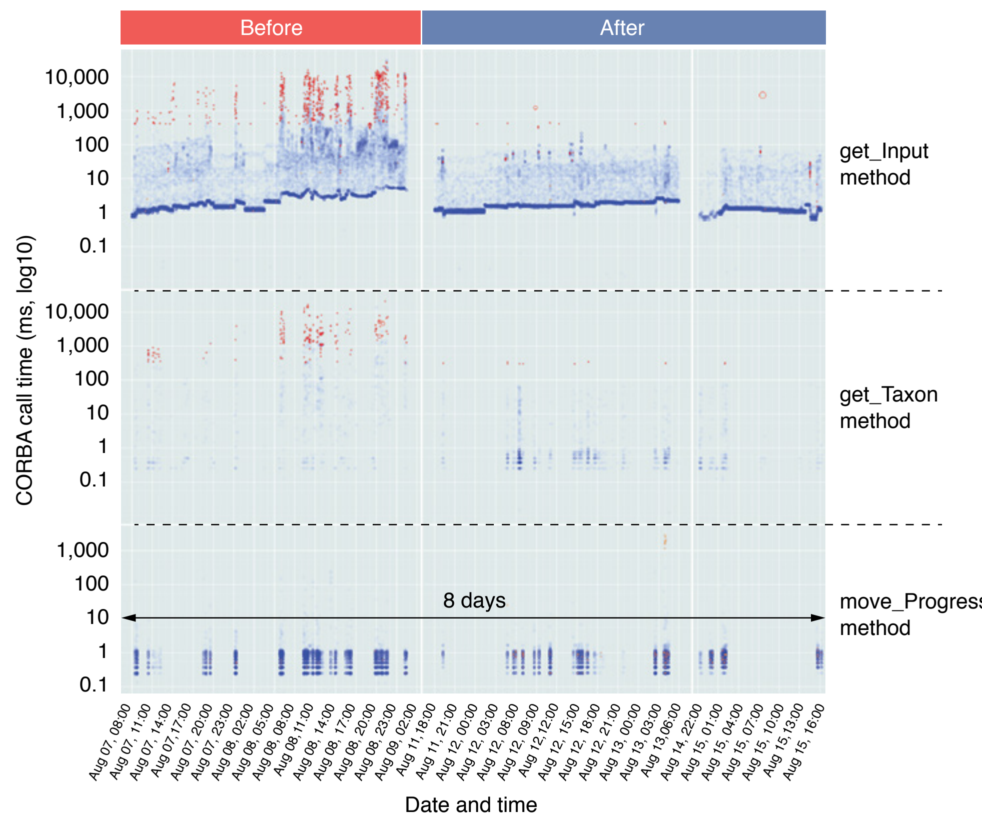

A call duration and retransmit diagram shows an off-normal pattern for the National Ignition Facility's Integrated Computer Control System. With the toolkit, the call durations and retransmissions were visualized, leading to resolution of the issue.

Creating a Path to Petascale-Class Uncertainty Quantification Studies


DAVID DOMYANCIC
DOMYANCIC1@LLNL.GOV

RAO NIMMAKAYALA
NIMMAKAYALA1@LLNL.GOV

UQ is a key technique for quantifying numerical error in multiphysics simulations. The LLNL UQ Pipeline is an automated end-to-end scientific workflow system that generates ensemble-of-models-based UQ studies. The pipeline utilizes Livermore Computing's (LC's) high performance computers to generate a set of simulations, or ensemble, whose input parameters have been perturbed according to several sampling methods. The pipeline has been used to produce the largest ensemble of climate models in the United States and to generate NIF ensemble studies using LC's terascale-class computing resources. Each study was typically composed of hundreds to thousands of simulations and had up to 30 perturbed parameters. Petascale-class UQ studies will consist of at least several thousand ensemble simulations, in which more than 30 parameters will need to be perturbed to better understand the predictive capability of modern multiphysics codes. Using LLNL's Sequoia supercomputer, we have executed ensemble-of-models studies and incorporated UQ methods to predict outcomes of untested experiments. These advancements directly contributed to making petascale class UQ studies a reality and will benefit many DOE research efforts such as inertial confinement fusion, climate change, weapons, and energy.

The UQ Pipeline's scheduler is an enabling technology that launches, executes, and monitors concurrent ensemble simulations automatically with minimal user involvement. In 2013, the Verification and Validation Team in the Laboratory's AX Division worked closely with LC computer scientists to port the scheduler to Sequoia and utilize the supercomputer's unique capabilities to execute UQ studies. To date, the pipeline has executed more than 8,000 ensemble simulations on Sequoia. Each simulation used 256 cores and executed up to 256 concurrent ensemble simulations, requiring up to 8 midplanes. The 8,000 simulations demonstrated that ensemble studies could be generated using Sequoia.

UQ analysis begins with generating an ensemble-of-models study for each tested and untested experiment. Statistical information obtained from the ensemble of models for experiments with observational data and uncertainties are used to make predictions about untested experiments. This methodology is called Multiexperiment Filtering (MEF), and in 2013, several MEF techniques were incorporated into the pipeline. One such method calculates the probability that each ensemble is consistent with the observed data. The probabilities are then used to predict, with uncertainty, the outcome of an untested experiment. While this method is well understood, its use is limited based on whether the results of the predictor experiments have a high level of agreement. The team developed a second MEF method called Statistical Filtering, a technique based loosely on evidence theory, which sums the results of the predictor experiments to make predictions.

Researchers used both MEF methods in 2013 to make uncertainty assessments of an untested experiment. Incorporating these analysis methods into the pipeline creates a fully automated UQ engine that generates ensemble-of-models calculations and then uses the MEF methods to make predictions with minimal user interaction. The progress made in 2013 to achieve petascale-class UQ studies also provides insight into how to execute these studies on future exascale systems.

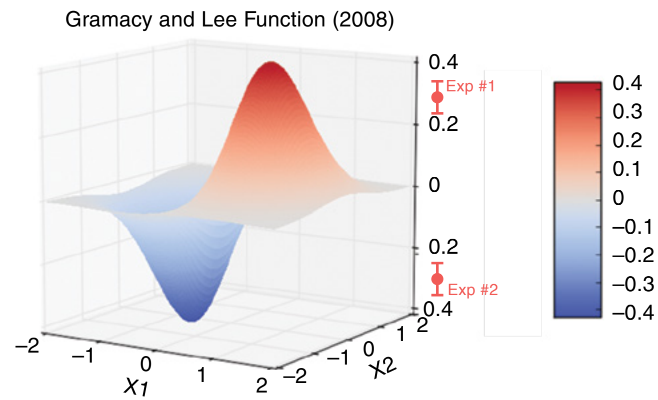

A surface plot of the Gramacy and Lee Function (2008) shows two experimental observations at –0.25 and 0.25 with standard deviations of 0.05 for each experiment (X1 and X2).

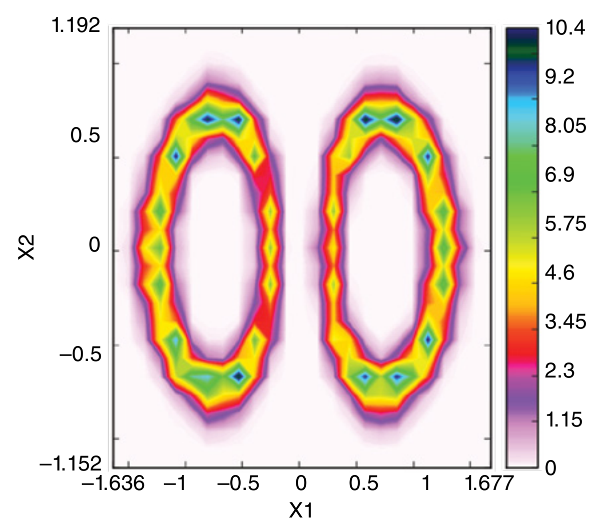

Statistical Filtering correctly identified both regions in parameter space that satisfy the experimental observations.

Providing Large-Scale Support to the Y-12 Complex

KIMBERLY FERRARI
FERRARI2@LLNL.GOV

More than 20 years after its initial deployment, Argus, the Laboratory's primary alarm and access control system, remains the gold standard for physical security systems at DOE sites. Argus is a comprehensive electronic security system designed to protect high-consequence assets through automated access controls, electronic monitoring, and a central security console. The LLNL-developed system operates 24 hours a day, 7 days a week and is flexible, adaptable, and capable of simultaneous monitoring of an entire site. More recent upgrades to Argus include the use of an Integrated Video Assessment System (IVAS) that provides surveillance capabilities using modern digital video technologies.

The Security and Protection (S&P) program is responsible for the Argus and IVAS systems. S&P personnel include professionals from the Computation and Engineering directorates, Security Organization, and Operations and Business Principal Directorate. In addition, S&P partners with Professional Project Services, Inc. (Pro2Serve), which provides critical infrastructure engineering services for national security. Last July, Computation experts and other LLNL personnel in the S&P program provided support to the Y-12 National Security Complex after three protesters unlawfully gained access to the site. The collaboration led to improvements in security processes and procedures.

The Argus security system monitors and controls entry into high-security areas at DOE sites, including Y-12, and simultaneously alerts and directs security forces to possible threats. On July 28, 2012, three individuals trespassed into the Y-12 National Security Complex. The intruders damaged government property and interfered with ongoing national defense operations but did not gain access to any nuclear materials. The Argus security system used at Y-12 alarmed appropriately during the breach.

Computation personnel provide normal line-item support for Y-12's Security Improvement Project, a multiyear effort to install a new Argus-based security framework at Y-12 that manages and integrates intrusion detection, alarm monitoring, and access to the site's control operations. After the July 28 incident, the S&P program provided on-site support, helping analyze the event logs and video regarding the breach. Personnel observed console operations and made recommendations for improving operational procedures. S&P and Computation experts also wrote a custom event log parser to assist Y-12 in determining false and nuisance alarm rates, and enable rates to be compared across the NNSA complex. This comparison allowed LLNL and Y-12 to work together to determine sensors in need of maintenance and tuning.

Y-12's overall system security design was also analyzed and improved as part of the post-incident effort. Many operations were moved from a centralized model to a distributed model, where the owners of the security elements are responsible for the security of their facilities—a model similar to LLNL and most other DOE sites. Additionally, a cutover was done to install an automated access control system into the protected area that integrates metal detection capabilities. This automation cuts down on the number of security personnel needed to manage people going into and out of the protected area. At peak times, 1,000 people enter or exit the Protected Area within an hour.

At the heart of Y-12 improvements is the Argus system, a capability that is continually enhanced by Livermore's computer programmers and engineers. Upgrades to the system are designed to keep abreast of security requirements and incorporate the latest advances in technology, helping to effectively protect the nation's high-consequence assets.

Argus contains an automated access control system for personnel and equipment, an electronic monitoring and alarm system for detecting intruders, and a central security console for monitoring events.

 

Livermore Loops Resurrected: Improving Compilers

RICH HORNUNG
HORNUNG1@LLNL.GOV

JEFF KEASLER
KEASLER1@LLNL.GOV

Theoretical peak performance of supercomputers grows dramatically with each new generation of hardware. However, complex computational science applications find it increasingly difficult to realize performance improvements at the same rate. Large applications, such as Advanced Simulation and Computing (ASC) multiphysics simulation codes, typically realize less than 10% of peak operation rates supported by current machines. This gap is widening as application performance increasingly relies on fine-grained parallelism, such as vectorization and threading.

A compiler translates software algorithms into machine instructions that are executed by a system's processors. The resulting efficiency of the executable file depends critically on how well a compiler does the translation and enables optimizations. To study this issue, LLNL computer scientists have developed the Livermore Compiler Analysis Loop Suite (LCALS). This proxy application focuses on compiler optimization quality and serves as a vehicle for co-design interactions between LLNL application developers and compiler and hardware vendors. It also provides a heretofore missing methodology to measure compiler quality as part of new machine procurements.

LCALS is a collection of loop kernels based partly on well-known "Livermore Loops" benchmarks developed in the 1970s and 1980s. These earlier suites, which included C and Fortran numerical kernels, were used to evaluate floating-point math performance on systems before porting full applications.

The primary intent of LCALS is to provide diverse compiler tests that span recurring idioms used in scientific software—for DOE applications in particular. LCALS contains kernels from older suites and adds more complex tests representative of modern applications. It is written in C++, the programming language for many LLNL ASC codes. It evaluates parallel execution features, such as single instruction, multiple data (SIMD) vectorization, Open Multi-Processing multithreading, and abstraction layers based on C++ template mechanisms.

In 2013, we used LCALS to establish ongoing dialogue with compiler vendors to address application performance and portability issues. The specific self-contained tests in LCALS allow a user to set up a controlled compilation and execution environment, run tests in that environment, and then isolate tests that deviate from expected efficiency. Tests that pinpoint problems are extracted and sent to vendors as individual bug reports, giving them a manageable set of small issues to address. As an example, LCALS uncovered SIMD optimization problems in an earlier version of the Intel C++ compiler. Using the data from the tests, Intel was able to resolve the detected issues, which ultimately will improve performance of LLNL application codes.

LCALS also provides a pathway for co-design activities. We are working with compiler vendors to define strategies to encapsulate nonportable software constructs in DOE applications and elsewhere. Data encapsulation strategies used in ASC codes often inhibit compiler optimizations or supply insufficient information to enable optimizations. In these cases, we can negotiate with vendors to define practical idioms that enable portable performance and meet both application and vendor requirements.

LCALS will be a marquee benchmark for the next large ASC system to be sited at LLNL in the 2017 time frame. This application of LCALS is the first time that a benchmark will be used to assess compiler capabilities as well as hardware performance. Through LCALS, we will work with vendors to address longstanding compiler performance issues for DOE applications, and these benefits could carry over to applications in the private sector.

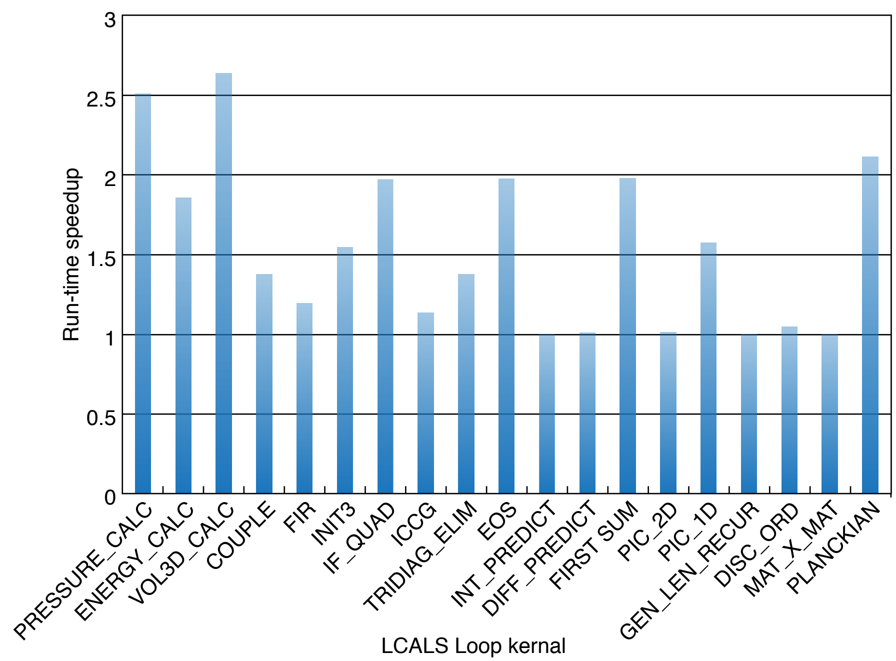

Run-time speedup improvement of part of the Livermore Compiler Analysis Loop Suite (LCALS) was observed using a new version of the Intel C++ compiler that resolved optimization issues LCALS uncovered in an earlier version.

Keeping Pace with the Data Deluge from Next-Generation Sequencing

JONATHAN ALLEN
ALLEN99@LLNL.GOV

Genome sequencing technology can recover billions of genetic fragments from an individual human clinical or environmental sample, making it possible to identify microbial pathogens without any prior knowledge of the organisms present. The approach known as metagenomic sequencing is helping reveal the complex relationship between the microbial world and human health for biosecurity applications. A major challenge is developing computational methods that accurately characterize samples with enough efficiency to scale the analysis to match the growth in sequencer use and output. The Livermore Metagenomic Analysis Toolkit (LMAT) is a new open source software tool designed to address this issue by offering precise, scalable metagenomic analysis.

Metagenomic classification involves comparing millions to billions of genetic fragments with unknown origin to a large database of reference genomes. Traditionally, each fragment is compared with every reference genome independently, and the matched genomes are combined to determine the originating organism. To avoid long run times, the reference database is limited in size and scope and usually stores only a representative set of bacterial strains. Organisms with larger genomes, such as fungi, are analyzed in rare cases but only when using sufficiently large computer clusters. LMAT was developed to address the fundamental algorithmic scaling limitations of existing metagenomic classification methods by using a custom reference database with a fast searchable index.

Previous approaches did not exploit important genetic relationships between the different reference organisms, a technique that enables fragment matching of multiple reference organisms. A key innovation of LMAT is that it can pre-compute the occurrence of each short (18–20 nucleotide) sequence across the entire database and store the evolutionarily conserved sequence patterns, which helps accelerate the classification process and improve the accuracy of the results. For example, within the custom database, a short sequence found in hundreds of bacterial strains can be recognized and replaced with a single species identifier to reduce costly, redundant comparisons. The current database tracks every available sequenced genome from viruses, bacteria, archaea, protozoa, fungi, and several variants of the human genome, tracking the conserved sequence patterns for approximately 25 billion short sequences.

The database uses up to 400GB of random access memory (RAM) and will continue to expand as larger eukaryotic genomes representing plants and animals are added. The growth in database size is managed using data structures optimized to support a multitiered memory hierarchy in combination with faster dynamic RAM and higher latency, lower cost persistent memory devices such as flash drives.

LMAT demonstrates the first metagenomic classification algorithm tailored to exploit the larger memory capacity offered on newer commercially available computers. The software has analyzed metagenomic data at a rate of one million nucleotides per second using a 1TB (of RAM), 40-core machine, ensuring that the speed of analysis matches that of the output from the fastest sequencers. LMAT can now run on a single compute node paired with individual sequencing machines to support true scaling of analysis and keep pace with the increase in sequencer use.

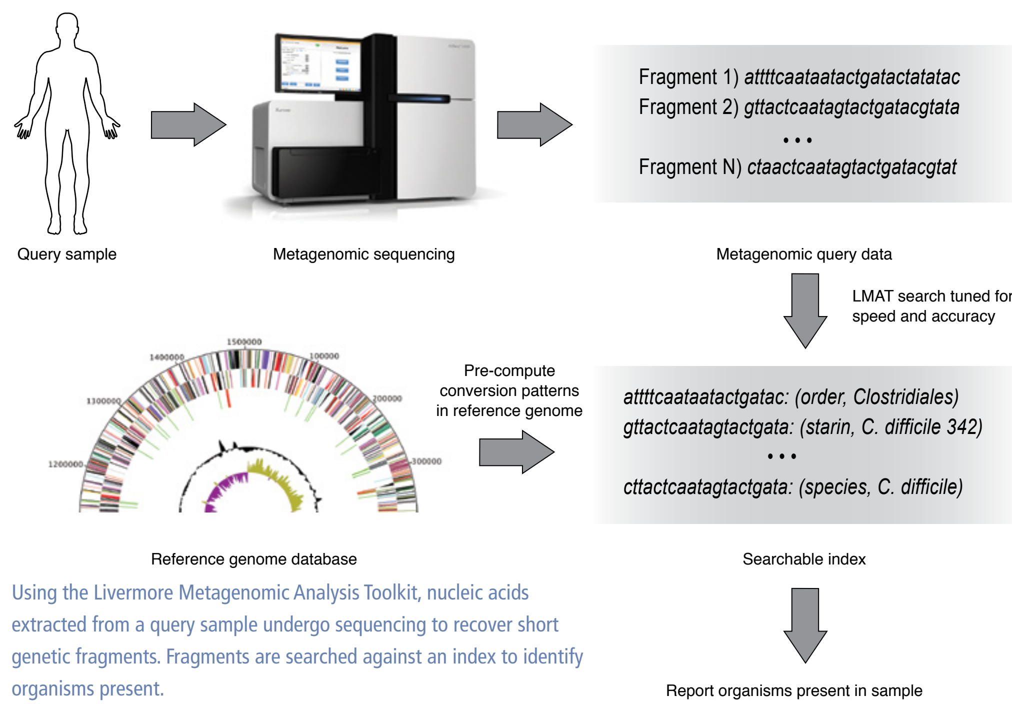

Using the Livermore Metagenomic Analysis Toolkit, nucleic acids extracted from a query sample undergo sequencing to recover short genetic fragments. Fragments are searched against an index to identify organisms present.

3 INSTITUTIONAL SERVICES

INSTITUTIONAL SERVICES FACILITATE THE ONE-LAB PRINCIPLE

ROBYNE TESLICH
TESLICH2@LLNL.GOV

LAWRENCE Livermore National Laboratory (LLNL) embraces a guiding principle, called "one Lab," wherein the institution pursues its missions and operates with a collective passion and shared responsibility, adopting operational efficiencies, common standards, and uniform business practices wherever possible.

The Institutional Services team supports the one-Lab vision for LLNL by providing services and common solutions that enable employee productivity and, as a result, further the overall missions of LLNL. Increasing demands from a collaborative and mobile workforce combined with budget constraints and more restrictive compliance requirements present a significant challenge to Computation personnel. Still, we remain committed to improving service delivery governance, increasing self-service capabilities, enabling our employees to work from various locations, and employing a risk-based approach to compliance and software development. We believe that centralized information technology (IT) services delivered according to best practices and common compliance processes, such as change management and software quality assurance (SQA), create efficiencies and improve product and service quality when implemented across the site.

In FY13, significant strides were made in IT to improve the efficient delivery of services and to provide additional capabilities to employees. Efficiencies were gained by consolidating support functions such as desktop and server security configuration management (e.g., patches and software distribution). In addition, tools were developed that enable end-user self-service for patch management, file recovery, and a knowledgebase containing support articles to assist users. The most significant progress has been in the area of mobility and secure data access for mobile users with the deployment of a new on-premise, cloud-based, file storage and sharing solution.

Balancing compliance and governance with agility is a continuous challenge. This year, we defined and implemented best-practice service delivery processes based on the Information Technology Infrastructure Library framework and SQA. We also implemented a common IT change management process across the IT infrastructure and operations scope, improving service change approvals and increasing customer awareness of service outages. SQA has become a part of LLNL software development culture, ensuring the appropriate amount of rigor is applied to development based on the intended use of the software.

Computation staff are committed to achieving balance between meeting compliance through risk-based decisions and mitigations and delivering leading-edge services and capabilities. Maintaining a competitive edge in the Department of Energy (DOE) mission space requires highly available, well-managed IT services for collaboration, communication and data management, and appropriate governance and quality rigor. Being responsive to mission needs while continuing to improve IT quality, security, mobility, and efficiency requires the sustained enhancement and consolidation of services. The LLNL mission, and ultimately LLNL sponsors, will benefit from the competitive advantage that the Institutional Services team offers to every LLNL employee.

myCloud Storage Service Enhances Employee Mobility, Collaboration, and Productivity

"The cloud" is an increasingly popular technology buzz phrase that covers various computing concepts. At its core, the cloud is a virtual location where users can put all their data, all their files, and even their software so they can access it from any computer or device, anywhere, anytime. In FY13, the LLNL Chief Information Officer (CIO) Program, with support from Computation staff, planned for and implemented a capability that allows Laboratory employees to synchronize their data in a secure, on-site private cloud and makes it accessible from their government-owned devices. This new service, called myCloud, enables LLNL employees to synchronize all or some of their unclassified data in real-time across all of their work desktops and mobile devices in the background seamlessly and securely while they work. myCloud provides secure access to the data from anywhere in the world without having to manually connect to the LLNL network. In addition, employees can choose to share files or entire folders with individuals or groups, without assistance from IT staff, making collaborations between employees easier than ever before.

In today's business world, employees demand and deserve seamless collaboration and mobility solutions that provide anywhere, anytime access to critical business information. The LLNL myCloud service helps accomplish that goal while maintaining the highest degree of information security to protect the information that is vital to the Laboratory mission.

myCloud offers Laboratory employees several features that combine to create a high-availability, on-site cloud data storage environment.

In the summer of 2013, the LLNL CIO Program began analyzing several large-scale Enterprise File Sync and Sharing (EFSS) solutions to fulfill the vision of an institution-wide storage service for employees to store their data files in a centrally managed repository. The CIO Program had spent the previous two years establishing an enterprise-class mobile device platform, so the time was right to seamlessly link an employee's desktop data to their mobile device. The chosen technology solution would need to provide a highly scalable, highly available infrastructure for more than 6,000 employees, be easy for the IT staff to deploy and manage, be simple for the employee to configure and use, and provide for full and automatic synchronization of an employee's data at all times, to all their devices, and do so while running transparently in the background. Also, the solution would need to be hosted on site in multiple LLNL data centers for redundancy yet be accessible from anywhere in the world with single sign-on access control (without RSA SecurID tokens or virtual private network authentication) so that employees would embrace its simplicity and convenience. Most importantly, the solution would need to pass strict cybersecurity requirements, including full National Institute of Standards and Technology-certified Federal Information Processing Standards 140-2 military-grade encryption for data in transit and data at rest for every device. After filtering the available products through these strict requirements and conducting a thorough product evaluation, EMC's Syncplicity EFSS was chosen as the backup and synchronization service on which to build LLNL's myCloud capability.

The Syncplicity orchestration is managed by EMC and is hosted in their data centers. Orchestration manages the communication, storage, and retrieval of data between clients and the compute nodes. LLNL has compute nodes on two unclassified networks for redundancy and data transfer efficiency. These compute nodes assemble and disassemble the data stored on the storage clusters. The back-end storage comprises redundant EMC Isilon storage clusters in geographically separate LLNL data centers with approximately 500TB of useable storage. This implementation model ensures that Laboratory data is only transferred between LLNL clients and LLNL compute nodes, and that all data is fully encrypted during transit.

myCloud's anywhere, anytime access to automatically synchronized data on multiple mobile and multiple desktop devices is an aspect of modern technology that today's sophisticated technology user has come to expect. One of the most unique features of myCloud is its simplicity of use, which belies the underlying complexity of technologies that empower the service. Once installed, the EMC Syncplicity software client that myCloud utilizes becomes fully integrated into the file system of a user's PC or Mac. Employees

CHUCK SELLERS
SELLERS4@LLNL.GOV

get the same user experience on either platform and can see the sync status of all their files directly from the file browser. Employees can then securely access their files, either offline or online, from their Laboratory computer or laptop as well as their mobile devices, regardless of where they originally created and saved them.

myCloud also makes it easier for LLNL scientists to collaborate on documents and files. There is no need to check files in and out; no emailing attachments to multiple people and then coordinating multiple response emails; and no need to manage different versions. Employees simply share a folder or file with their colleagues directly from where the file is stored. Changes made after sharing a file will be shared as well, so everyone's version is always up to date. Permissions can be set to allow others to collaborate (edit), or files can be shared as read-only. The Syncplicity client automatically handles version control and syncing the shared folders and files across all devices with the latest version.

myCloud also greatly reduces the risk of data loss by replicating data files to multiple computers, laptops, and mobile devices. If a desktop computer malfunctions, the user's data will be on another device. If the user does not have an additional device, once the damaged device is repaired or replaced, he or she can simply reload the data from the master copy stored in the LLNL data center. If an older version of a file is needed, the user can recover a previous version or a deleted file without the assistance of LLNL IT staff. If the user is in another part of the Laboratory without any of their devices and needs a specific file, he or she can use a web browser to access the data.

With the combination of both a scalable and highly available mobile device service platform with an equally world-class EFSS storage service, the LLNL workforce now has increased mobility, productivity, and security, which provides them with a technological and competitive edge.

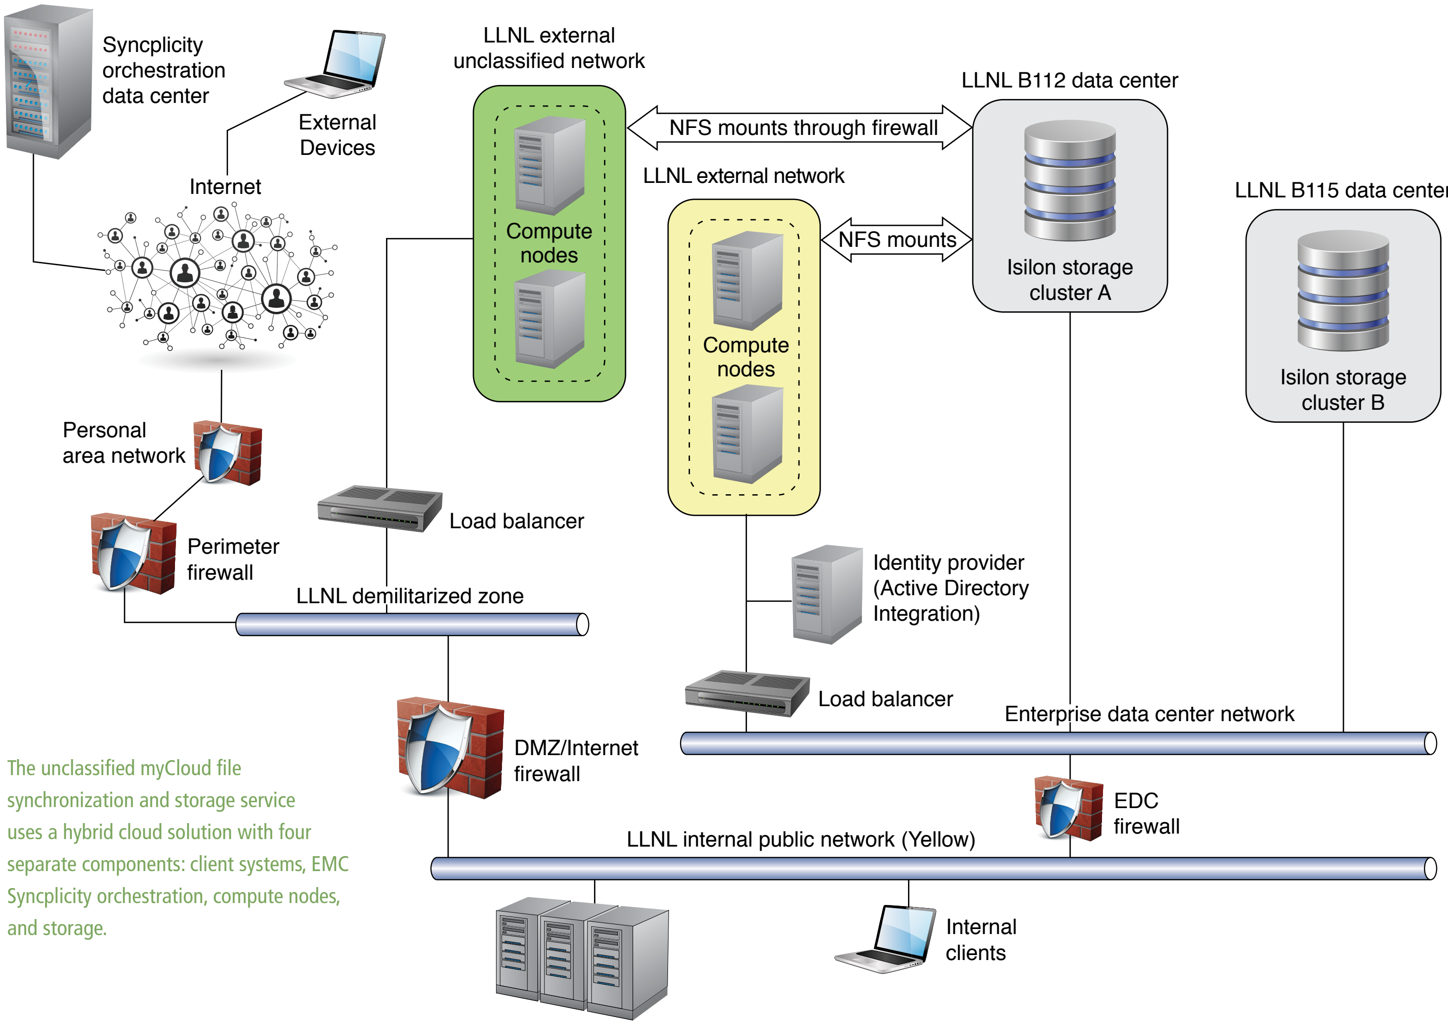

The unclassified myCloud file synchronization and storage service uses a hybrid cloud solution with four separate components: client systems, EMC Syncplicity orchestration, compute nodes, and storage.

LANDesk Infrastructure Consolidation Creates Efficiencies

GREG HERWEG
HERWEG1@LLNL.GOV

The Computation team that supports the CIO Program continues to find ways to improve IT and create efficiencies at the Laboratory. These improvements are necessary to conserve critical investments in LLNL infrastructure, allow IT staff to focus on more challenging hands-on issues, and foster a healthier work environment with enhanced productivity for computer users. Several IT infrastructure management systems operate behind the scenes to maintain a safe computing environment that is both robust and complies with cybersecurity requirements. Historically, the responsibility for operating some of these IT management systems has fallen on many different programs. The Computation team's multiyear objective is to eliminate redundancy and combine IT services where possible, while continuing to support programmatic missions with minimal disruption in day-to-day operations.

One of the major accomplishments this year in the area of IT service consolidation was combining multiple instances of LLNL's patching system (LANDesk) into a single LANDesk infrastructure. LANDesk is LLNL's primary tool for patching, deploying software packages, and otherwise managing several thousand unclassified Windows operating system-based computers.

LANDesk was initially deployed in 2006 in a very distributed way, with each organization responsible for its own set of systems. Although this strategy was adequate, it was not efficient. The update workflow required a good amount of manual labor by the IT staff in each organization. In addition, patches were being installed on systems at various times and days, thus hindering the productivity of users. So, in 2013, the CIO Program combined all the separate installations of the LANDesk tool into a single instance that is centrally managed, thereby significantly decreasing the burden on LLNL's IT staff. The consolidation has also made more efficient the day-to-day operation and maintenance of the core Windows patching and software deployment tool and has improved the process for incorporating software changes into the IT ecosystem at LLNL. Overall, the experience of updating or replacing software is now smoother and more consistent for all end-users across the enterprise.

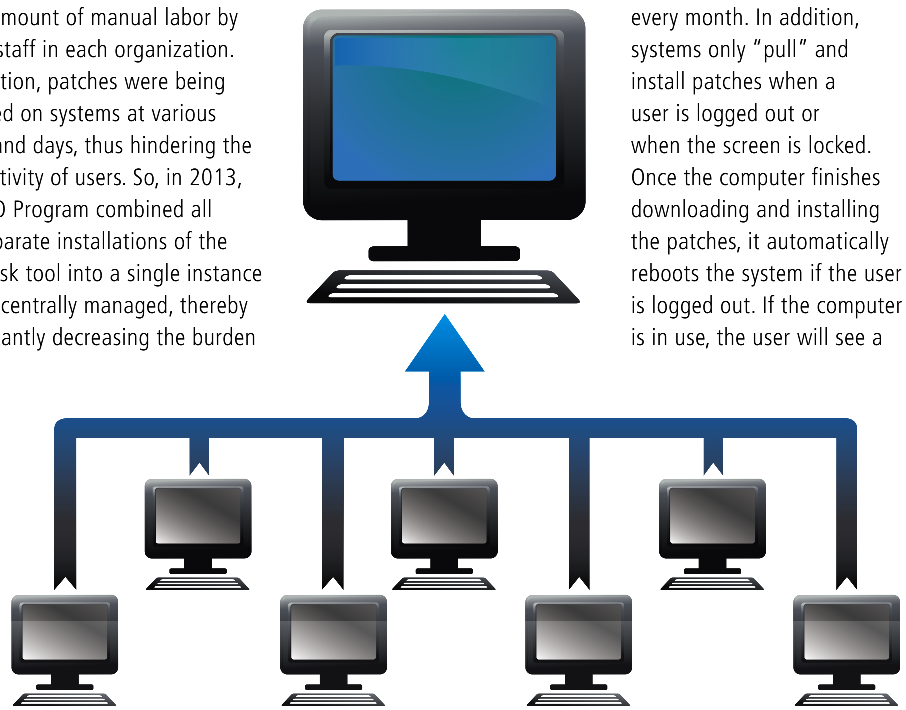

New policies on the central LANDesk server help minimize the impact of patching on users' workdays. The LANDesk patching process now occurs on a specific day of every month. In addition, systems only "pull" and install patches when a user is logged out or when the screen is locked. Once the computer finishes downloading and installing the patches, it automatically reboots the system if the user is logged out. If the computer is in use, the user will see a reboot prompt that will time out and auto-reboot after 72 hours.

LLNL IT administrators use various software tools to manage the many thousands of Windows, Macintosh, and Linux computer systems at the Laboratory. Some tools, like LANDesk, have already been migrated to a more unified approach. However, other tools, including critical desktop/laptop maintenance tools such as backup software and hardware, are currently deployed and managed in a distributed, decentralized manner. The next phase of our consolidation project will likely focus on the numerous backup systems that exist across the site.

The CIO Program has laid the groundwork and established philosophies for providing increased automation for day-to-day desktop/laptop computer operations. Having the economy of scale that consolidation brings will allow us to deliver more end-user self-help tools and resolve issues more quickly, resulting in higher user satisfaction and increased productivity for LLNL employees.

Improving the information technology ecosystem at the Laboratory involves a multiyear effort to eliminate redundancy and combine services.

Software Quality Innovations Balance Agility and Discipline

GREGORY POPE
POPE12@LLNL.GOV

The research software that is developed at LLNL provides scientific insight at much faster speeds and greater precision than software developed only a few years ago. The dynamic nature of our software requires an SQA approach that balances agility and discipline and ensures the appropriate development rigor based on the software's intended usage. In 2013, LLNL software developers and software quality engineers (SQEs) implemented process improvements and collaborated with other Department of Energy (DOE) laboratories to broaden the scope of the Laboratory's SQA approach.

Championing software quality has become part of LLNL's culture. In 2013, LLNL SQEs and scientific software developers made several quality innovations that improve specific phases of the development and testing cycle.

For example, debugging software that runs across thousands of processors is daunting due to the massive volume of data the software encompasses. It is much easier for humans to spot imperfections in visual representations of data than to identify an incorrect number in a vast series. This realization was the impetus for a team of software developers and SQEs to bundle together VisIt, an LLNL open-source viewer, and Rogue Wave's TotalView, a parallel debugger. The result is a tool prototype that enables visual representations, differentiations, and comparisons of data during the debugging phase.

Another challenge that was addressed this year was the increasing complexity of building new scientific codes using existing component packages. While code reuse is desirable, the complexity of the interdependencies of the packages and the order in which they must be built can be overwhelming for programmers. To simplify the task, a tool called MixDown was created, which helps programmers analyze a dependency graph to determine the proper ordering of packages as well as which packages can be configured and built concurrently. MixDown is a DOE Advanced Scientific Computing Research-funded scientific software project that has reduced total development effort and simplified the task of building software with multiple dependencies.

Another daunting task programmers face is determining the correctness of the equation-of-state data used to represent materials for simulations. This important data is hosted on a Linux platform with a graphic user interface (GUI) application for data viewing. By enhancing the GUI with a test automation framework (froglogic's Squish) that runs under Linux, programmers can now exhaustively test the data through the GUI. Squish has reduced the time spent on GUI testing software releases while increasing the quality of the application.

Software agility is supported by the continuous integration and testing of codes as soon as they are developed. To support continuous integration, LLNL developers and SQEs implemented the open-source Jenkins tool on Advanced Simulation and Computing projects. Jenkins detects changes to the source code in the repository, builds and tests the new code on multiple platforms, and provides test reports nightly or multiple times daily.

LLNL's risk-based SQA approach is highly regarded within the DOE community. In 2013, LLNL supported the C++ codes and provided our SQA tools and capabilities to six other DOE laboratories. One such shared tool, static analysis, finds structural code issues that are often missed by compilers and traditional testing methods. In addition, DOE's Office of Nuclear Energy uses LLNL's approach on their Nuclear Energy Advanced Modeling and Simulation project and has requested an abbreviated version as a model for all programs in the Office of Nuclear Energy to consider.

A talented team of professionals was instrumental in advancing the cause of software quality at the Laboratory in FY13, including (standing, from left) Tamara Dahlgren, Evelyn Chen, Jeff Keasler, Esteban Pauli, Christopher White, Cyrus Harrison, (seated, from left) Thomas Epperly, Stephanie Dempsey, and Bill Oliver.

New Change Management Process Enhances IT Service Delivery

REA SIMPSON
SIMPSON6@LLNL.GOV

IT is a critical, underlying component of nearly every LLNL mission; thus, a program's success depends on high-quality, reliable IT services. Because the majority of IT service outages result from changes to IT infrastructure, products, or systems, Computation staff supporting the CIO Program led the implemention of a program-wide change management process. The goals of the project were to establish standardized methods for handling all IT changes, to minimize the impact of such changes on the Laboratory's business, and to facilitate a faster rate of change. Since implementing the formal change management process, IT staff are able to respond more quickly to new business needs. Also, the number and severity of change-related service outages have been greatly reduced. This project reached the top 3 in the IT Excellence Awards: Project of the Year competition at the IT Service Management Conference and Exhibition in February 2013.

The CIO Program consists of six functional areas: Institutional Services, System Management Services, Business Applications, Data Center, Cyber Security, and Telecommunications. A few of the areas already had change management processes in place, but process maturity varied and some areas had no process at all. The first stage of the project involved analyzing the few existing processes and leveraging, rather than wholesale replacing, the aspects that were working well. Our deliverables included a documented change management process with defined roles and responsibilities, a detailed communication plan for changes based on the level of impact, and defined maintenance windows.

One major challenge was to create a change process that struck the appropriate balance between required documentation, review and approvals, and impact to the efficiency of the organization. We defined roles and responsibilities that included a Change Process Owner (this champion of the effort ensures the process is being carried out as defined and also initiates process improvements) and a Change Manager (the person responsible for carrying out the process). Services were then categorized based on their criticality to the business. The level of service criticality combined with the risk of the change determines the level of approval required before the change can be scheduled. The highest level of approval occurs at the Change Advisory Board, which includes representatives from all involved areas of IT as well as impacted customers.

Communication is another critical component of change management. Our team developed a detailed communication plan that identified how, when, and to whom communication should occur based on the level of impact of the change. We also developed templates so that all communications have a common appearance and include common elements, regardless of which area is initiating the communication.

Prior to the existence of a formal change management process, business-impacting IT changes were not restricted to specific maintenance windows. The ad hoc scheduling of changes made it difficult for customers to plan their work. It also resulted in change "collisions" within the IT infrastructure. Now, standard maintenance windows have been established, and the work that occurs during each window is managed within the change process. The timing of major work (e.g., complete network or storage outages) is scheduled and communicated at the beginning of the year, giving customers enough advanced notice to plan accordingly.

The new change management process, including roles and process flow, was automated within LLNL's IT Service Management tool. This automation allows us to efficiently execute the process and generate metrics that track and monitor progress.

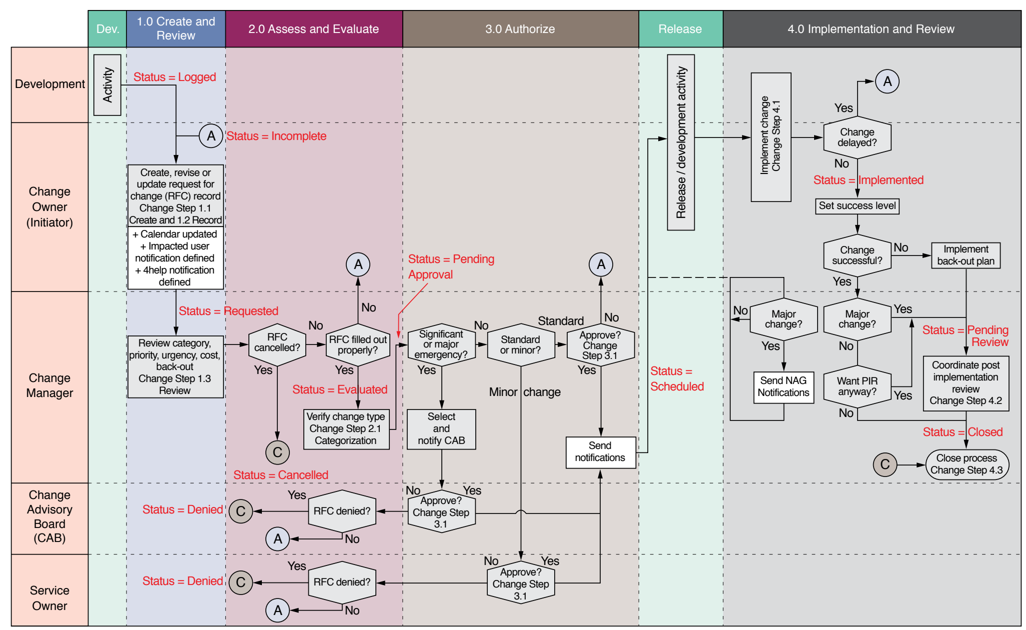

The change management process has four phases: Create and Review, Assess and Evaluate, Authorize, and Implementation and Review. "Release activities" to plan and prepare for the change occur between the Authorize and Implementation-and-Review phases.

4 APPENDICES

ACADEMIC OUTREACH

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
Allen University	Abdollah Rabieh	Joint research	Malware intelligence harvesting for greater cyber defense	NNSA	Matt Myrick
Arizona State University	Stephen Johnson	Joint research	Peptide microarrays	DHS	Tom Slezak
Boston University	Jonathan Appavoo	Collaboration	Operating systems and runtime	ASCR	Maya Gokhale
Brigham Young University	Bryan Morse	Joint research	Mosaics and super-resolution of unmanned aerial vehicle-acquired video using locally adaptive warping	LDRD	Mark Duchaineau and Jon Cohen
California Institute of Technology	Dale Pullin	Collaboration	Turbulence modeling	ACSR Base	Bill Henshaw
California Polytechnic State University, San Luis Obispo	David Clague	Joint research	Low-cost microarrays	LDRD	Tom Slezak
California Polytechnic State University, San Luis Obispo	Ignatios Vakalis	Joint research	Cybersecurity research: joint proposals	SMS	Celeste Matarazzo
Cambridge University	Nikos Nikiforkis	Joint research	Simulation and modeling using Overture	ASCR Base	Bill Henshaw
Carnegie Mellon University	Christos Faloutsos	Joint research	Mining large, time-evolving data for cyber domains; joint proposals	—	Celeste Matarazzo
Carnegie Mellon University	Franz Franchetti	Joint research	Performance optimization of fast Fourier transform on Blue Gene/Q	ASC	Martin Schulz
Carnegie Mellon University	Christos Faloutsos	Subcontract	Network change detection and attribution	LDRD	Brian Gallagher
Chalmers University of Technology	Sally McKee	Collaboration	Leveraging OpenAnalysis for alias analysis within ROSE	ASC	Dan Quinlan
Clark Atlanta University	Roy George	Joint research	Malware intelligence harvesting for greater cyber defense	NNSA	Matt Myrick
Clark Atlanta University	Peter Molnar	Joint research	Cybersecurity research; joint proposals		Celeste Matarazzo
Colorado State University	Donald Estep	Subcontract	A posteriori error analysis for hydrodynamic systems	LDRD	Carol Woodward
Colorado State University	Stephen Guzik	Subcontract	A node-level programming model framework for exascale computing	LDRD	Chunhua Liao
Colorado State University	Michelle Strout and Sanjay Rajopadhye	Collaboration	Program analysis	ASCR	Dan Quinlan
Columbia University	Mark Adams	Joint research	FASTMath: Frameworks, Algorithms, and Scalable Technologies for Mathematics	ASCR SciDAC	Lori Diachin
Columbia University	Ian Lipkin	Joint research	Viral discovery and microarrays	DTRA	Tom Slezak
Cornell University	Ken Birman	Joint research	Evaluation of scalable cloud computing technologies for use in Department of Energy systems and applications	ASCR	Greg Bronevetsky
Darmstadt University of Technology	Christian Bischof	Joint research	OpenMP performance tools	ASC	Martin Schulz

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
Dresden University of Technology	Wolfgang Nagel	Joint research	Improved analysis of message-passing interface traces and performance measurement infrastructures	ASC	Martin Schulz
Dresden University of Technology	Wolfgang Nagel	Joint research	Semantic debugging of message-passing interface applications; trace-based performance analysis	ASCR, ASC, LDRD	Bronis de Supinski
Dresden University of Technology	Wolfgang Nagel and Andreas Knuepfer	Joint research	Message-passing interface trace analysis	ASC	Martin Schulz
ETH Zürich	Thorsten Hoefler	Joint research	Message-passing interface forum and advanced message-passing interface usage, performance and power modeling	ASC	Martin Schulz
Georgetown University	Heidi Wachs	Joint research	Cybersecurity research: joint proposals	—	Celeste Matarazzo
Georgia Institute of Technology	Raheem Beyah	Joint research	Cybersecurity research: joint proposals	—	Celeste Matarazzo
Georgia Institute of Technology	Richard Fujimoto	Subcontract	Research in reverse computation	LDRD	David Jefferson
Georgia Institute of Technology	Jarek Rossignac	Subcontract	Compact streamable mesh formats	ASCR	Peter Lindstrom
Georgia Institute of Technology	Richard Vuduc	Subcontract	Compiler support for reverse computation	ASCR	Dan Quinlan
Georgia Institute of Technology	Dan Campbell and Mark Richards	Collaboration	Data-intensive applications	WFO/DARPA	Maya Gokhale
Imperial College	Paul Kelly and José Gabriel de Figueiredo Coutinho	Collaboration	Field-programmable gate arrays research	ASCR	Dan Quinlan
Imperial College	Anthony Lim	Collaboration	Time-dependent density functional theory Qbox simulations	—	Erik Draeger
Indiana University	Jeremiah Wilcock	Joint research	Binary analysis	ASCR	Dan Quinlan
Johns Hopkins University	Allan Boyles	Collaboration	Seismoacoustic modeling for defense-related efforts	DOE	Shawn Larsen
Kyushu University, Japan	Koji Inoue	Joint research	Energy- and power-aware high performance computing	ASC, ASCR	Martin Schulz
Louisiana State University	Lu Peng, Lide Duan, and Sui Chen	Joint research	Characterizing the propagation of soft faults through numeric applications	ASCR	Greg Bronevetsky
Ludwig Maximilians University of Munich	Dieter Kranzlmueller	Joint research	Detecting communication patterns to optimize applications	ASCR, ASC	Bronis de Supinski
Ludwig Maximilians University of Munich	Dieter Kranzlmueller	Joint research	Message-passing interface tool infrastructure and performance analysis, power- and energy-aware high performance computing	ASC, ASCR	Martin Schulz
Naval Medical Research Center	Vish Mokashi	Joint research	Microbial forensics	DTRA	Tom Slezak
Norfolk State University	Aftab Ahmad and Jonathan Graham	Joint research	Cybersecurity research: joint proposals	SMS	Celeste Matarazzo

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
North Carolina Agricultural and Technical State University	Gerry Dozier	Joint research	Cybersecurity research: joint proposals	—	Celeste Matarazzo
North Carolina Agricultural and Technical State University	Gerry Dozier	Joint research	Malware intelligence harvesting for greater cyber defense	NNSA	Matt Myrick
Northern Arizona University	Paul Keim	Joint research	Microbial forensics	DTRA	Tom Slezak
Ohio State University	P. Sadayappan and Christophe Alias	Collaboration	Optimizing compiler program analysis	ASCR	Dan Quinlan
Ohio University	Yusu Wang	Joint research	Analysis and visualization of high-dimensional function	LDRD	Timo Bremer
Pennsylvania State University	Ludmil Zikatonov	Subcontract	Multilevel methods for graph Laplacians and piece-wise constant approximations	LDRD	Panayot Vassilevski
Pennsylvania State University	Ludmil Zikatonov and Robert Scheichl	Subcontract	Construction of coarse-scale numerical models based on energy minimization and applications to upscaling techniques	ASCR Base	Panayot Vassilevski
Pennsylvania State University	Jinchao Xu and James Brannick	Subcontract	Multigrid methods for systems of partial differential equations	ASCR	Robert Falgout
Pennsylvania State University	Jinchao Xu	Collaboration	Multigrid solvers theory	—	Tzanio Kolev
Polytechnic University of Puerto Rico	Alfredo Cruz	Joint research	Cybersecurity research: joint proposals	—	Celeste Matarazzo
Portland State	Jay Gopalakrishnan	Collaboration	Discontinuous Petrov-Galerkin methods	—	Tzanio Kolev
Purdue University	Saurabh Bagchi	Joint research	Anomaly detection and tracking in high performance computing	ASC	Martin Schulz
Purdue University	Saurabh Bagchi	Subcontract	Root cause analysis of faults in parallel systems	ASCR	Greg Bronevetsky
Purdue University	Saurabh Bagchi	Joint research	Statistical debugging tools, fault tolerance, scalable checkpointing	ASCR, ASC, LDRD	Bronis de Supinski
Purdue University	Zhiqiang Cai	Summer faculty	A posteriori error estimates for partial differential equations	ASC	Robert Falgout
Purdue University	Jennifer Neville	Joint research	Hypothesis tests for dynamic networks	LDRD	Brian Gallagher
Purdue University	Mithuna Thottethodi	Joint research	Optimized node-mapping techniques	ASC	Martin Schulz
Purdue University	Dongbin Xiu	Subcontract	Intrusive and nonintrusive polynomial chaos methods for the quantification of uncertainties in multiphysics models	ASC	Charles Tong
Queens University of Belfast	Dimitris Nikolopoulos	Joint research	Power optimization for hybrid codes, epidemiology simulation	ASC	Martin Schulz
Rensselaer Polytechnic Institute	Don Schwendeman	Subcontract	Development of numerical methods for mathematical models of high-speed reactive and nonreactive flow	ASCR Base	Bill Henshaw

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
Rensselaer Polytechnic Institute	Mark Shephard and Onkar Sahni	Joint research	FASTMath: Frameworks, Algorithms, and Scalable Technologies for Mathematics	ASCR SciDAC	Lori Diachin
Rice University	John Mellor-Crummey	Joint research	Performance analysis, standardization for OpenMP	ASC, ASCR	Martin Schulz
Rice University	John Mellor-Crummey	Joint research	Sustained Performance, Energy, and Resilience Institute	ASCR SciDAC	Bronis de Supinski
Rice University	Vivek Sarkar	Subcontract	Data abstractions for portable high performance computing performance	LDRD	James McGraw
Rice University	John Mellor-Crummey, Keith Cooper, and Vivek Sarkar	Collaboration	Use of ROSE for compiler optimizations	ASCR	Dan Quinlan
Rice University	Vivek Sarkar, Jisheng Zhao, Vincent Cave, and Micheal Burke	Joint research	Development of a static single assignment-based dataflow compiler framework for ROSE	ASCR	Greg Bronevetsky
Rice University	Vivek Sarkar, Jisheng Zhao, Vincent Cave, and Micheal Burke	Subcontract	Message-passing interface producer–consumer program analyses for ROSE	ASCR	Greg Bronevetsky
Rochester Institute of Technology	Kara Maki	Collaboration	Droplet flows	—	Bill Henshaw
Royal Institute of Technology, Sweden	Heinz-Otto Kreiss	Consultant	Adaptive methods for partial differential equations	ASCR Base	Anders Petersson
Rutgers University	Tina Eliassi-Rad	Subcontract	Cyber situational awareness through host and network analysis	LDRD	Celeste Matarazzo
Rutgers University	Tina Eliassi-Rad	Subcontract	Network similarity	LDRD	Brian Gallagher
RWTH Aachen University	Matthias Mueller	Joint research	Message-passing interface correctness checking	ASC	Martin Schulz
RWTH Aachen University	Felix Wolf and Matthias Muller	Joint research	Message-passing interface performance analysis tools	ASCR, ASC, LDRD	Bronis de Supinski
RWTH Aachen University	Felix Wolf	Joint research	Parallel performance analysis	ASC, ASCR	Martin Schulz
RWTH Aachen University	Matthias Muller	Joint research	Message-passing interface correctness checking	ASC	Martin Schulz
Southern Methodist University	Thomas Hagstrom	Joint research	High-order structure grid methods for wave propagation on complex unbounded domains	ASCR Base	Bill Henshaw
Southern Methodist University	Dan Reynolds	Joint research/subcontract	FASTMath: Frameworks, Algorithms, and Scalable Technologies for Mathematics	ASCR SciDAC	Lori Diachin
Southern Methodist University	Dan Reynolds	Subcontract	New time-integration methods and support for multiscale solution methods in the LLNL SUNDIALS software suite	ASCR SciDAC	Carol Woodward
Stanford University	Juan Alonso	Subcontract	New modes of laser lethality	LDRD	Kambiz Salari
Stanford University	Sanjiva Lele	Subcontract	Development of a nonequilibrium wall model for the compressible flow solver CharLESX	DOE	Kambiz Salari

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
Stanford University	Olav Lindtjorn	Collaboration	Reverse-time seismic imaging for hydrocarbon exploration	CRADA	Shawn Larsen
Stanford University	Subhasish Mitra	Joint research	Quantifying the accuracy of fault injection tools that operate at difference system abstraction levels	ASCR	Greg Bronevetsky
Stanford University	Gianluca Iaccarino	ASC Predictive Science Academic Alliance Program Center	Predictive Simulations of Particle-Laden Turbulence in a Radiation Environment	ASC	Kambiz Salari and Bronis de Supinski
Technical University of Denmark	Sven Karlsson	Joint research	Scalable debugging	ASC	Martin Schulz and Dong Anh
Technical University of Denmark	Robert Read	Collaboration	Water waves and wave energy generation	—	Bill Henshaw
Technical University of Munich	Arndt Bode	Joint research	Exascale computing	ASC	Martin Schulz
Technical University of Vienna	Markus Schordan	Collaboration	Compiler construction	ASCR	Dan Quinlan
Texas A&M University	Nancy Amato	Joint research	Load balance optimizations	ASC	Martin Schulz
Texas A&M University	Nancy Amato	Collaboration, Lawrence Scholar Program	Novel mechanisms to understand and improve load balance in message-passing interface applications	UCOP	Bronis de Supinski
Texas A&M University	Nancy Amato	Collaboration, Lawrence Scholar Program	Parallel graph algorithms	UCOP	Maya Gokhale
Texas A&M University	Yalchin Efendiev	Joint research	Algebraic multigrid and Bayesian uncertainty quantification for Darcy and Brinkman problems	ASCR	Panayot Vassilevski
Texas A&M University	Jean-Luc Guermond and Boian Popov	Collaboration	Artificial viscosity for Lagrangian hydrodynamics	—	Tzanio Kolev
Texas A&M University	Jim Morel	ASC Predictive Science Academic Alliance Program Center	Center for Exascale Radiation Transport	ASC	Teresa Bailey and Todd Gamblin
Texas A&M University	Bjarne Stroustrup and Lawrence Rauchwerger	Joint research	Compiler construction and parallel optimizations	ASCR	Dan Quinlan
Texas State University	Byron Gao	Collaboration	Search user interfaces; clustering	LDRD	David Buttler
Tufts University	Scott MacLachlan	Joint research	Parallel multigrid in time	ASCR	Robert Falgout
UC Berkeley	Domagoj Babic	Joint research	Formal models of communication protocols	ASCR	Greg Bronevetsky
UC Berkeley	Doug Dreger	Collaboration	Earthquake hazard	IGPP	Shawn Larsen
UC Berkeley	Edgar Solomonik	Collaboration	Communication-avoiding parallel linear algebra algorithms for use in Qbox	—	Erik Draeger

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
UC Berkeley	Per-Olof Persson	Collaboration	Space-time Discontinuous Galerkin algorithms	—	Tzanio Kolev
UC Davis	François Gygi	Collaboration	General Qbox development, new algorithms	—	Erik Draeger
UC Davis	Bernd Hamann	Joint research	Analysis and visualization of performance data	UC Fee	Timo Bremer
UC Davis	Bernd Hamann	Joint research	Performance analysis and visualization	ASC	Martin Schulz
UC Davis	Ken Joy	Joint research	Visualization and Analytics Center for Enabling Technologies	ASCR SciDAC	Jon Cohen
UC Davis	Ken Joy	Subcontract	Improving accuracy and efficiency of 3D aerial video	DoD	Jon Cohen
UC Davis	Ken Joy	Collaboration, Lawrence Scholar Program	Discrete multimaterial interface reconstruction for volume fraction data	UCOP	Jon Cohen
UC Davis	Ken Joy	Collaboration, Lawrence Scholar Program	Video-processing research for the VidCharts and Persistics projects	UCOP	Jon Cohen
UC Davis	Louise Kellogg	Joint research	Topological analysis of geological data	LDRD	Timo Bremer
UC Davis	Kwan-Liu Ma	Subcontract	Interactive tomographic reconstruction of aerial motion imagery	DoD	Mark Duchaineau
UC Davis	Nelson Max	Joint research	3D from wide-area aerial video	DOE Nonproliferation	Mark Duchaineau
UC Davis	Zhendong Su	Subcontract	ROSE support for rewrapping macro calls	ASCR	Dan Quinlan
UC Davis	Matt Bishop and Sean Peisert	Joint research	Cybersecurity research, joint proposals, cyber defenders	—	Celeste Matarazzo
UC Davis	Sean Peisert	Subcontract	Network resilience	LDRD	Brian Gallagher
UC Riverside	Michalis Faloutsos	Joint research	Cybersecurity research: joint proposals	—	Celeste Matarazzo
UC San Diego	Randy Bank	Subcontract	Solvers for large sparse systems of linear equations	ASCR	Robert Falgout
UC San Diego	Laura Carrington	Joint research	Sustained Performance, Energy, and Resilience Institute	ASCR SciDAC	Bronis de Supinski
UC San Diego	Laura Carrington and Scott Baden	Joint research	Performance modeling	ASC	Martin Schulz
UC San Diego	Falko Kuester	Joint research	3D from video	DoD	Mark Duchaineau
UC San Diego	Tai Cheung	Collaboration	Relationship between technology and national security in China	UC Fee	Dona Crawford
UC San Diego	Falko Kuester	Joint research	Large-scale atomistic simulation visualization	ASC	Mark Duchaineau
UC San Diego	Laura Carrington	Collaboration	Data-intensive architectures	UC Fee	Maya Gokhale
UC San Diego	Steve Swanson	Collaboration	Persistent memory emulator	—	Maya Gokhale

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
UC San Diego	Erik Gartzke and Jon Lindsey	Joint research	Cybersecurity research: joint proposals	—	Celeste Matarazzo
UC San Diego	Yuri Bazilevs	Collaboration	Nonuniform rational basis spline (NURBS) discretizations and isogeometric analysis	—	Tzanio Kolev
UC San Diego, Scripps Institution of Oceanography	Julie McClean	Collaboration	Ultra-high-resolution coupled climate simulations	BER	Art Mirin
UC Santa Cruz	Steve Kang	Collaboration	Persistent memory devices	UC Fee	Maya Gokhale
UC Santa Cruz	Carlos Maltzahn	Collaboration, Lawrence Scholar Program	Semantic file systems	LDRD	Maya Gokhale
United States Army Medical Research Unit, Kenya	John Waitumbi	Joint research	Pathogen diagnostics	—	Tom Slezak
University of Arizona	David Lowenthal	Joint research	Power-aware computing for message-passing interface programs; scalable performance models	ASCR, ASC	Bronis de Supinski
University of Arizona	David Lowenthal	Joint research	Power optimization and modeling	ASC, ASCR	Martin Schulz
University of British Columbia	Carl Olivier-Gooch	Subcontract	Aggressive mesh optimization	ASCR SciDAC	Lori Diachin
University of British Columbia	Carl Olivier-Gooch	Subcontract	FASTMath: Frameworks, Algorithms, and Scalable Technologies for Mathematics/Mesquite	ASCR SciDAC/ASCR Base	Lori Diachin
University of Colorado	Ken Jansen	Joint research	FASTMath: Frameworks, Algorithms, and Scalable Technologies for Mathematics	ASCR SciDAC	Lori Diachin
University of Colorado	Tom Manteuffel	Joint research	Solution methods for transport problems	ASC	Peter Brown
University of Colorado	Steve McCormick and Tom Manteuffel	Subcontract	Adaptive algebraic multigrid for graph mining problems	LDRD	Van Emden Henson
University of Colorado	Steve McCormick, Tom Manteuffel, John Ruge, and Marian Brezina	Subcontract	Error estimators for uncertainty quantification, adaptive mesh refinement, solvers for Stochastic partial differential equations, parallel adaptive algebraic multigrid/smoothed aggregation, and parallel solution of systems of partial differential equations	ASC	Robert Falgout
University of Colorado, Denver	Andrew Knyazev	Subcontract	Improving efficiency and accuracy of O(N) solvers in finite-difference density functional theory calculations	LDRD	Jean-Luc Fattebert
University of Delaware	Richard Braun	Collaboration	Models of the eye	ASCR Base	Bill Henshaw
University of Delaware	John Cavazos	Subcontract	ROSE compiler project	ASCR	Dan Quinlan and Chunhua Liao
University of Florida	S. Bala Balachandar	ASC Predictive Science Academic Alliance Program Center	Center for Compressible Multiphase Turbulence	ASC	Sam Schofield and Maya Gokhale

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
University of Illinois at Urbana-Champaign	William Gropp	Joint research	Message-passing interface, hybrid programming models	ASCR, ASC, LDRD	Bronis de Supinski
University of Illinois at Urbana-Champaign	William Gropp	Joint research	Optimization for algebraic multigrid	ASC	Martin Schulz
University of Illinois at Urbana-Champaign	William Gropp	ASC Predictive Science Academic Alliance Program Center	Center for Exascale Simulation of Plasma-Coupled Combustion	ASC	Alan Kuhl and Rob Neely
University of Illinois at Urbana-Champaign	Laxmikant Kale	Subcontract	Node-mapping optimizations for high performance computing systems	ASC, LDRD	Abhinav Bhatele
University of Illinois at Urbana-Champaign	Laxmikant Kale and Esteban Meneses	Joint research	Scalable check pointing and message logging for fault-tolerant high performance computing systems	ASCR	Greg Bronevetsky
University of Illinois at Urbana-Champaign	Rakesh Kumar and Joseph Sloan	Joint research	Algorithm-specific fault tolerance for sparse linear algebra applications	ASCR	Greg Bronevetsky
University of Illinois and IBM	Hormozd Gahvari, William Gropp, Luke Olson, and Kirk Jordan	Collaboration	Modeling algebraic multigrid performance on multicore architectures	LDRD	Ulrike Yang
University of Karlsruhe	Wolfgang Karl	Joint research	Hardware transactional memory	ASC	Martin Schulz
University of Louisville	Yongsheng Lian	Collaboration	Micro-air vehicles	—	Bill Henshaw
University of Maryland	Jeffrey Hollingsworth	Joint research	Autotuning and tool infrastructures	ASCR	Martin Schulz
University of Maryland	Jeffrey Hollingsworth	Joint research	Sustained Performance, Energy, and Resilience Institute	ASCR SciDAC	Bronis de Supinski
University of Massachusetts, Amherst	Andrew McCallum	Joint research	Cross-language topic models	LDRD	David Buttler
University of Nevada, Reno	John Louie	Collaboration	Seismic modeling in the basin and range region	DOE	Shawn Larsen
University of New Mexico	Dorian Arnold	Joint research	Scalable tool infrastructures	ASC, ASCR	Bronis de Supinski
University of New Mexico	Dorian Arnold	Joint research	Tool infrastructures	ASC	Greg Lee
University of North Carolina	Robert Fowler	Joint research	Sustained Performance, Energy, and Resilience Institute	ASCR SciDAC	Bronis de Supinski
University of North Carolina	Jan Prins	Joint research	Efficient OpenMP runtimes for tasking	ASC, LDRD	Bronis de Supinski
University of North Carolina	Jan Prins	Joint research	OpenMP task scheduling	ASC	Martin Schulz
University of Notre Dame	Samuel Paolucci	ASC Predictive Science Academic Alliance Program Center	Center for Shock Wave-Processing of Advanced Reactive Materials	ASC	Nathan Barton and Brian Spears

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
University of Ontario	Isaac Tamblyn	Collaboration	Time-dependent density functional theory Qbox simulations	—	Erik Draeger
University of San Francisco	Jeff Buckwalters	Joint research	Performance modeling	ASC	Martin Schulz
University of Southern California	Gerard Medioni	Subcontract	Activity analysis in wide-area overhead video	DOE Nonproliferation	Jon Cohen
University of Southern California	Robert Lucas and Jacqueline Chame	Joint research	Sustained Performance, Energy, and Resilience Institute	ASCR SciDAC	Bronis de Supinski
University of Tennessee	Jack Dongarra	Joint research	Empirical tuning	ASCR	Dan Quinlan
University of Tennessee	Jack Dongarra and Shirley Moore	Joint research	Sustained Performance, Energy, and Resilience Institute	ASCR SciDAC	Bronis de Supinski
University of Texas, San Antonio	Shirley Moore	Joint research	Sustained Performance, Energy, and Resilience Institute	ASCR SciDAC	Bronis de Supinski
University of Texas Medical Branch	Yuriy Fofanov	Joint research	Genomic algorithms	DTRA	Tom Slezak
University of Turabo	Jeffrey Duffany	Joint research	Cybersecurity research: analysis and defense of large-scale smart meter networks; joint proposals	SMS	Celeste Matarazzo
University of Utah	Ganesh Gopalakrishnan	Subcontract	Compiler analysis of message-passing interface applications	ASCR	Greg Bronevetsky
University of Utah	Ganesh Gopalakrishnan	Collaboration	Message-passing interface optimizations	ASCR	Dan Quinlan
University of Utah	Ganesh Gopalakrishnan	Joint research	Identification and targeted elimination of nondeterminism	ASC	Dong Ahn
University of Utah	Ganesh Gopalakrishnan and Mary Hall	Joint research	Semantic debugging of message-passing interface applications, and Sustained Performance, Energy, and Resilience Institute	ASCR, ASC, LDRD	Bronis de Supinski
University of Utah	Valerio Pascucci	Joint research	Performance analysis and visualization	ASC, ASCR	Martin Schulz
University of Utah	Valerio Pascucci	Subcontract	Performance analysis and visualization	LDRD	Timo Bremer
University of Utah	Philip Smith	ASC Predictive Science Academic Alliance Program Center	Carbon-Capture Multidisciplinary Simulation Center	ASC	Bert Still and Greg Burton
University of Utah	Chris Johnson, Valerio Pascucci, Chuck Hansen, Claudio Silva, Lee Myers, Allen Sanderson, and Steve Parker	Joint research	Visualization and Analytics Center for Enabling Technologies	ASCR SciDAC	Mark Duchaineau
University of Washington	Carl Ebeling and Scott Hauck	Collaboration	Coarse-grain processor architectures	—	Maya Gokhale
University of Waterloo	Hans de Sterck	Subcontract	Numerical methods for large-scale data factorization	LDRD	Van Emden Henson

UNIVERSITY	FACULTY	ACTIVITY TYPE	TOPIC	FUNDING SOURCE	LLNL CONTACT
University of Wisconsin	Bart Miller	Joint research	Scalable debugging	ASCR, ASC, LDRD	Bronis de Supinski
University of Wisconsin	Karu Sankaralingam	Joint research	Fault tolerant computing models for high performance computing	ASC	Martin Schulz
University of Wisconsin	Karu Sankaralingam	Joint research	Resilient computing	ASCR, ASC	Bronis de Supinski
University of Wisconsin	Bart Miller and Ben Liblit	Joint research	Performance tools and tool infrastructures	ASCR, ASC	Martin Schulz
Utah State University	Renée Bryce and Steena Monteiro	Joint research	Statistical modeling of data-driven applications	ASCR, Lawrence Scholar Program	Greg Bronevetsky
Virginia Institute of Technology	Kirk Cameron	Joint research	Power-aware computing for hybrid systems	ASCR, ASC	Bronis de Supinski
Virginia Institute of Technology	Wu-chun Feng	Joint research	Hybrid computing programming models, power-aware computing	ASCR, ASC, LDRD	Bronis de Supinski
Virginia Institute of Technology	Madhav Marathe	Joint research	Mathematical and computational foundations of network sciences	SMS	Celeste Matarazzo
Virginia Institute of Technology	Madhav Marathe	Subcontract	Research in network sciences	SMS	Evi Dube and James Brase
Virginia Institute of Technology	Madhav Marathe	Joint research	Epidemiology simulation at scale	ASC	Martin Schulz
Voorhees College	Tim Kentopp	Joint research	Malware intelligence harvesting for greater cyber defense	NNSA	Matt Myrick

PUBLICATIONS

JOURNAL ARTICLES AND CONFERENCE PROCEEDINGS

Adams, M. P., “Provably Optimal Parallel Transport Sweeps on Regular Grids,” *International Conference on Mathematics and Computational Methods Applied to Nuclear Science & Engineering (M&C 2013)*, (American Nuclear Society), (LLNL-CONF-618792).

Ames, S., M. Gokhale, and C. Maltzahn, “QMDS: A File System Metadata Management Service Supporting a Graph Data Model-Based Query Language,” *International Journal of Parallel, Emergent and Distributed Systems* **28** (2), 159–183 (LLNL-JRNL-518611).

Ames, S., et al., “Scalable Metagenomic Taxonomy Classification Using a Reference Genome Database,” *Bioinformatics* **29** (18), 2253–2260.

Antipa, N. A., et al., “Automated ICF Capsule Characterization Using Confocal Surface Profilometry,” *Fusion Sci. Technol.* **63** (2), 151–159.

Antipa, N. A., et al., “Automated ICF Capsule Characterization Using Confocal Surface Profilometry,” *Target Fabrication Meeting*, (Fusion Science and Technology), **63** (2), 151.

Awwal, A. A. S., et al., “Image Processing and Control of a Programmable Spatial Light Modulator for Spatial Beam Shaping,” Conference on High Power Lasers for Fusion Research II,” *Proceedings of the International Society for Optics and Photonics*, (SPIE Int. Soc. Optical Engineering), **8602**, (LLNL-CONF-613592).

Banks, J. W. and T. D. Aslam, “Richardson Extrapolation for Linearly Degenerate Discontinuities,” *J. Sci. Comput.* **57** (1), 1–18 (LLNL-JRNL-558251).

Banks, J. W., J. A. F. Hittinger, J. M. Connors, and C. S. Woodward, “A Posteriori Error Estimation Via Nonlinear Error Transport with Application to Shallow Water,” *8th International Conference on Scientific Computing and Applications*, (Contemp. Math., Amer. Mathematical Soc.), **586**, 35–42 (LLNL-JRNL-562712).

Banks, J. W., W. D. Henshaw, and B. Sjogreen, “A Stable FSI Algorithm for Light Rigid Bodies in Compressible Flow,” *J. Comput. Phys.* **245**, 399–430 (LLNL-JRNL-558232).

Bhatia, H., G. Norgard, V. Pascucci, and P.-T. Bremer, “Comments on the Meshless Helmholtz–Hodge Decomposition,” *IEEE T. Vis. Comput. Gr.* **19** (3), 527–528.

Bhatia, H., G. Norgard, V. Pascucci, and P.-T. Bremer, “The Helmholtz–Hodge Decomposition: A Survey,” *IEEE T. Vis. Comput. Gr.* **19** (8), 1386–1404.

Bihari, B. L., “Transactional Memory for Unstructured Mesh Simulations,” *J. Sci. Comput.* **54** (2–3), 311–332 (LLNL-JRNL-484451).

Borucki, M. K., et al., “The Role of Viral Population Diversity in Adaptation of Bovine Coronavirus to New Host Environments,” *PLoS ONE* **8** (1).

Carman, L., “The Effect of Material Purity on the Optical and Scintillation Properties of Solution-Grown Trans-Stilbene Crystals,” *J. Cryst. Growth* **368**, 56–61.

Carnes, B., et al., “Science at LLNL with IBM Blue Gene/Q,” *IBM J. Res. Dev.* **57** (1–2).

Chatrchyan, S., et al., “Evidence for Associated Production of a Single Top Quark and W Boson in pp Collisions at s = 7 TeV,” *Phys. Rev. Lett.* **110** (2).

Chaudhury, S., et al., “Rapid Countermeasure Discovery against *Francisella tularensis* Based on a Metabolic Network Reconstruction,” *PLoS ONE* **8** (5).

Chen, M. J., et al., “Surrogate-Based Optimization of Hydraulic Fracturing in Preexisting Fracture Networks,” *Comput. Geosci.-UK* **58**, 69–79 (LLNL-JRNL-628332).

Chen, X., B. Ng, Y. W. Sun, and C. Tong, “A Flexible Uncertainty Quantification Method for Linearly Coupled Multiphysics Systems,” *J. Comput. Phys.* **248**, 383–401.

Chen-Harris, H., et al., “Ultra-Deep Mutant Spectrum Profiling: Improving Sequencing Accuracy Using Overlapping Read Pairs,” *BMC Genomics* **14**.

Clark, D. S., et al., “Detailed Implosion Modeling of Deuterium–Tritium Layered Experiments on the National Ignition Facility,” *Phys. Plasmas* **20** (5).

Cleveland, M. A., T. A. Brunner, N. A. Gentile, and J. A. Keasler, “Obtaining Identical Results with Double Precision Global Accuracy on Different Numbers of Processors in Parallel Particle Monte Carlo Simulations,” *J. Comput. Phys.* **251**, 223–236.

Connors, J. M., J. W. Banks, J. A. Hittinger, and C.S. Woodward, “A Method to Calculate Numerical Errors Using Adjoint Error Estimation for Linear Advection,” *SIAM J. Numer. Anal.* **51** (2), 894–926.

Correa, C.D. and P. Lindstrom, “The Mutual Information Diagram for Uncertainty Visualization,” *International Journal for Uncertainty Quantification* **3** (3), 187–201 (LLNL-JRNL-558392).

Day, R., et al., “Understanding the General Packing Rearrangements Required for Successful Template-Based Modeling of Protein Structure from a CASP Experiment,” *Comput. Biol. Chem.* **42**, 40–48.

DeAngelis, K. M., et al., “Metagenomes of Tropical Soil-Derived Anaerobic Switchgrass-Adapted Consortia with and without Iron,” *Stand. Genomic Sci.* **7** (3), 382–398.

D’haeseleer, P., et al., “Proteogenomic Analysis of a Thermophilic Bacterial Consortium Adapted to Deconstruct Switchgrass,” *PLoS ONE* **8** (7).

Dobrev, V. A., T. E. Ellis, T. V. Kolev, and R. N. Rieben, “High-Order Curvilinear Finite Elements for Axisymmetric Lagrangian Hydrodynamics,” *Comput. Fluids* **83**, 58–69.

Dorf, M. A., et al., “Numerical Modeling of Geodesic Acoustic Mode Relaxation in a Tokamak Edge,” *Nucl. Fusion* **53** (6).

Dorf, M. A., et al., “Simulation of Neoclassical Transport with the Continuum Gyrokinetic Code COGENT,” *Phys. Plasmas* **20** (1).

Eder, D. C., A. C. Fisher, A. E. Koniges, and N. D. Masters, “Modeling Debris and Shrapnel Generation in Inertial Confinement Fusion Experiments,” *Nucl. Fusion* **53** (11), 113037.

Edwards, M. J., et al., “Progress Toward Ignition on the National Ignition Facility,” *Phys. Plasmas* **20** (7).

Fang, F., et al., “The Independent Set Perturbation Method for Efficient Computation of Sensitivities with Applications to Data Assimilation and a Finite Element Shallow Water Model,” *Comput. Fluids* **76**, 33-49.

Feigenbaum, E., R. A. Sacks, K. P. McCandless, and B. J. MacGowan, “Algorithm for Fourier Propagation Through the Near-Focal Region,” *Appl. Optics* **52** (20), 5030–5035.

Friedley, A., et al., “Ownership Passing: Efficient Distributed Memory Programming on Multicore Systems,” *ACM Sigplan Notices* **48** (8), 177–186.

Gaffney, J. A., D. Clark, V. Sonnad, and S. B. Libby, “Bayesian Inference of Inaccuracies in Radiation Transport Physics from Inertial Confinement Fusion Experiments,” *High Energ. Dens. Phys.* **9** (3), 457–461.

Gaffney, J. A., D. Clark, V. Sonnad, and S. B. Libby, “Development of a Bayesian Method for the Analysis of Inertial Confinement Fusion Experiments on the NIF,” *Nucl. Fusion* **53** (7).

Gahvari, H., et al., “Modeling the Performance of an Algebraic Multigrid Cycle Using Hybrid MPI/OpenMP,” *41st International Conference on Parallel Processing* (LLNL-CONF-533431).

Gahvari, H., et al., “Performance Modeling of Algebraic Multigrid on Blue Gene/Q: Lessons Learned,” *3rd International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computing Systems (PMBS12)* (LLNL-CONF-580692).

Gao, B. J., et al., “User-Centric Organization of Search Results,” *IEEE Internet Comput.* **17** (3), 52–59.

Gardner, S. N. and C. J. Jaing, “Bioinformatics for Microbial Genotyping of Equine Encephalitis Viruses, Orthopoxviruses, and Hantaviruses,” *J. Virol. Meth.* **193** (1), 112–120.

Gardner, S. N., et al., “Optimizing SNP Microarray Probe Design for High Accuracy Microbial Genotyping,” *J. Microbiol. Meth.* **94** (3), 303–310.

Gel, A., et al., “Applying Uncertainty Quantification to Multiphase Flow Computational Fluid Dynamics,” *Powder Technol.* **242**, 27–39.

Gerber, S., et al., “Morse–Smale Regression,” *J. Comput. Graph. Stat.* **22** (1), 193–214.

Goehner, J. D., et al., “LIBI: A Framework for Bootstrapping Extreme-Scale Software Systems,” *Parallel Comput.* **39** (3), 167–176.

Grim, G. P., et al., “Nuclear Imaging of the Fuel Assembly in Ignition Experiments,” *Phys. Plasmas* **20** (5).

Gurung, T., M. Luffel, P. Lindstrom, and J. Rossignac, “Zipper: A Compact Connectivity Data Structure for Triangle Meshes,” *Comput. Aided Design* **45** (2), 262–269 (LLNL-JRNL-503831).

Haan, S. W., et al., “NIF Ignition Campaign Target Performance and Requirements: Status May 2012,” *Fusion Sci. Technol.* **63** (2), 67–75.

Harrison, C. D. and H. Krishnan, "Python's Role in VisIt," *11th Annual Scientific Computing with Python Conference (SciPy 2012)*.

Harrison, C., et al., "Efficient Dynamic Derived Field Generation on Many-Core Architectures Using Python," *Workshop on Python for High Performance and Scientific Computing (PyHPC)*, (ACM).

Hittinger, J. A. F. and J. W. Banks, "Block-Structured Adaptive Mesh Refinement Algorithms for Vlasov Simulation," *J. Comput. Phys.* **241**, 118–140 (LLNL-JRNL-515291).

Hoen, A. G., S. N. Gardner, and J. H. Moore, "Identification of SNPs Associated with Variola Virus Virulence," *Biodata Min* **6**.

Howell, D., et al., "First Determination of the 8Li Valence Neutron Asymptotic Normalization Coefficient Using the 7Li(8Li, 7Li) 8Li Reaction," *Phys. Rev. C* **88** (2), (LLNL-JRNL-622632).

Howell, L., "Grafting AMR Capability Into the Diffusion Package of an Existing ALE Code," *Nuclear Explosive Code Development Conference* (LLNL-CONF-609323).

Hu, X. Z., P. S. Vassilevski, and J. C. Xu, "Comparative Convergence Analysis of Nonlinear Amli-Cycle Multigrid," *SIAM J. Numer. Anal.* **51** (2), 1349–1369.

Islam, T., et al., "MCRENGINE: A Scalable Checkpointing System Using Data-Aware Aggregation and Compression," *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC12* (LLNL-CONF-554251).

Karlin, I., J. McGraw, J. Keasler, and C. Still, "Tuning the LULESH Mini-App for Current and Future Hardware," *Nuclear Explosive Code Development Conference* (LLNL-CONF-610032).

Karlin, I., et al., "Exploring Traditional and Emerging Parallel Programming Models Using a Proxy Application," *27th IEEE International Parallel and Distributed Processing Symposium* (LLNL-CONF-586774).

Keyes, D. E., et al., "Multiphysics Simulations: Challenges and Opportunities," *Int. J. High Perform. C.* **27** (1), 4–83.

Kirkwood, R. K., et al., "A Review of Laser–Plasma Interaction Physics of Indirect-Drive Fusion," *Plasma Phys. Contr. F.* **55** (10).

Kline, J. L., et al., "Hohlraum Energetics Scaling to 520 TW on the National Ignition Facility," *Phys. Plasmas* **20** (5).

Laguna, I., et al., "Probabilistic Diagnosis of Performance Faults in Large-Scale Parallel Applications," *21st International Conference on Parallel Architecture and Compilation Techniques (PACT)*.

Lamarque, J. F., et al., "Multimodel Mean Nitrogen and Sulfur Deposition from the Atmospheric Chemistry and Climate Model Intercomparison Project (ACCMIP): Evaluation of Historical and Projected Future Changes," *Atmos. Chem. Phys.* **13** (16), 7997–8018 (LLNL-JRNL-639257).

Lang, D. M., A. T. Zemla, and C. L. E. Zhou, "Highly Similar Structural Frames Link the Template Tunnel and NTP Entry Tunnel to the Exterior Surface in RNA-Dependent RNA Polymerases," *Nucleic Acids Res.* **41** (3), 1464–1482 (LLNL-JRNL-602932).

Lani, A., B. Sjogreen, H. C. Yee, and W. D. Henshaw, "Variable High-Order Multiblock Overlapping Grid Methods for Mixed Steady and Unsteady Multiscale Viscous Flows, Part II: Hypersonic Nonequilibrium Flows," *Commun. Comput. Phys.* **13** (2), 583–602.

Li, D., et al., "Strategies for Energy-Efficient Resource Management of Hybrid Programming Models," *IEEE T. Parall. Distrib. Sys.* **24** (1), 144–157 (UCRL-JRNL-521391).

Li, J., et al., "Assessing Parameter Importance of the Common Land Model Based on Qualitative and Quantitative Sensitivity Analysis," *Hydrol. Earth Syst. Sci.* **17** (8), 3279–3293.

Lucas, D. D., et al., "Failure Analysis of Parameter-Induced Simulation Crashes in Climate Models," *Geosci. Model Dev.* **6** (4), 1157–1171.

Mauger, S. A., et al., "Self-Assembly of Selective Interfaces in Organic Photovoltaics," *Adv. Funct. Mater.* **23** (15), 1935–1946.

Meezan, N. B., et al., "X-Ray Driven Implosions at Ignition Relevant Velocities on the National Ignition Facility," *Phys. Plasmas* **20** (5).

Norgard, G. and P.-T. Bremer, "Ridge-Valley Graphs: Combinatorial Ridge Detection Using Jacobi Sets," *Comput. Aided Geom. D* **30** (6), 597–608.

Norgard, G. and P.-T. Bremer, "Robust Computation of Morse–Smale Complexes of Bilinear Functions," *Comput. Aided Geom. D* **30** (6), 577–587.

Rajanna, C., et al., "A Strain of *Yersinia pestis* with a Mutator Phenotype from the Republic of Georgia," *FEMS Microbiol. Lett.* **343** (2), 113–120.

Richards, D. F., et al., "Toward Real-Time Simulation of Cardiac Electrophysiology in a Human Heart at High Resolution," *Comput. Method. Biomec.* **16** (7), 802–805.

Rountree, B., et al., "Parallelizing Heavyweight Debugging Tools with MPIecho," *Parallel Comput.* **39** (3), 156–166.

Santer, B. D., et al., "Identifying Human Influences on Atmospheric Temperature," *P. Natl. Acad. Sci. USA* **110** (1), 26–33.

Santos, E., et al., "UV-CDAT: Analyzing Climate Datasets from a User's Perspective," *Comput. Sci. Eng.* **15** (1), 94–103.

Sato, K., et al., "Design and Modeling of a Nonblocking Checkpointing System," *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC12*.

Sjogreen, B. and J. W. Banks, "Stability of Finite Difference Discretizations of Multiphysics Interface Conditions," *Commun. Comput. Phys.* **13** (2), 386–410.

Stevens, K., K. Kegelmeyer, P. Andrzejewski, and D. Buttler, "Exploring Topic Coherence Over Many Models and Many Topics," *Empirical Methods in Natural Language Processing.*

Stevens, K., T. Huang, and D. Buttler, "The C-Cat Wordnet Package: An Open Source Package for Modifying and Applying Wordnet," *Global Wordnet Conference 2012* (LLNL-CONF-499791).

Su, C.-Y., et al., "Model-Based, Memory-Centric Performance and Power Optimization on NUMA Multiprocessors," *International Symposium on Workload Characterization* (LLNL-CONF-589272).

Sun, Y., et al., "Global Sampling for Integrating Physics-Specific Subsystems and Quantifying Uncertainties of CO_2 Geological Sequestration," *Int. J. Greenh. Gas Con*. **12**, 108–123.

Szafaryn, L. G., T. Gamblin, B. R. de Supinski, and K. Skadron, "Trellis: Portability Across Architectures with a High-Level Framework," *J. Parallel Distr. Com*. **73** (10), 1400–1413.

Tan, Y., et al., "PREPARE: Predictive Performance Anomaly Prevention for Virtualized Cloud Systems," *2012 IEEE 32nd International Conference on Distributed Computing Systems*, **32**, 285–294.

Weiss, K., F. Iuricich, R. Fellegara, and L. De Floriani, "A Primal/Dual Representation for Discrete Morse Complexes on Tetrahedral Meshes," *Comput. Graph. Forum* **32** (3), 361–370.

Williams, D. N., et al., "Ultrascale Visualization of Climate Data," *Computer* **46** (9), 68–76.

Winjum, B. J., et al., "Kinetic Simulations of the Self-Focusing and Dissipation of Finite-Width Electron Plasma Waves," *Phys. Rev. Lett.* **111** (10).

Xu, H. M., et al., "Simulation of Explosion Ground Motions Using a Hydrodynamic-to-Elastic Coupling Approach in Three Dimensions," *B. Seismol. Soc. Am.* **103** (3), 1629–1639 (LLNL-JRNL-557566).

Zapata, J. C., et al., "An Attenuated Lassa Vaccine in SIV-Infected Rhesus Macaques Does Not Persist or Cause Arenavirus Disease But Does Elicit Lassa Virus-Specific Immunity," *Virol. J.* **10**.

Zhan, C. S., X. M. Song, J. Xia, C. Tong, "An Efficient Integrated Approach for Global Sensitivity Analysis of Hydrological Model Parameters," *Environ. Modeling Softw.* **41**, 39–52.

Zhang, W., et al., "Castro: A New Compressible Astrophysical Solver III Multigroup Radiation Hydrodynamics," *Astrophys. J. Suppl. S.* **204** (1), 1–27.

INDUSTRIAL COLLABORATIONS

COMPANY	TOPIC	LLNL CONTACT
Adaptive Computing Enterprises, Inc.	Moab workload manager	Don Lipari
Affymetrix	Microarray evaluation	Tom Slezak and Shea Gardner
AMD	FastForward processor and memory project	Robin Goldstone, Martin Schulz, and Kathryn Mohror
AMD	Power and energy	Barry Rountree
AMD	Two-level memory, processing in memory	Maya Gokhale, Scott Lloyd, and Brian Van Essen
AOSense, Inc.	Gravity gradiometry in the detection of anomalous mass distribution in vehicles	Vijay Sonnad
Argo Navis Technologies	Automated cache performance analysis and optimization in Open|SpeedShop	Kathryn Mohror and Barry Rountree
Arista Networks	Low-latency Ethernet networks	Matt Leininger
ARM	Processing in memory	Maya Gokhale and Scott Lloyd
Battelle	Terrorism risk assessments	Amy Waters and Lisa Belk
Catalyst Technologies	Exploring system software and applications algorithms	Matt Leininger
Cisco, NetApp	Big data cloud computing	Robin Goldstone
Cisco Systems, Dell, DataDirect Networks, Intel, NetApp, Mellanox Technolgoies, QLogic, Red Hat, Oracle, and Supermicro	Hyperion collaboration	Matt Leininger
Commissariat à l'Énergie Atomique	Resource management software	Don Lipari, Jim Garlick, Mark Grondona, and Dong Ahn
Cray	Scalable capacity clusters	Matt Leininger and Trent D'Hooge
Cray	Exploring the Chapel programming language using LULESH	Abhinav Bhatele
Cymer	Extreme ultraviolet simulation and analysis	Jeff Wolf and others
DataDirect Networks	RAID 6 research and development for I/O systems	Mark Gary
Dell Computers	Scalable capacity clusters	Matt Leininger and Trent D'Hooge
Electric Power Research Institute	High performance computing for power grid simulations	Liang Min
Électricité de France	Aeroacoustics	Bill Henshaw
Energy Exemplar	PLEXOS power grid optimization	Jeff Wolf
EOG Resources	Seismic processing	Shawn Larsen
ExxonMobil	Cooperative Research and Development Agreement: computational mathematics	John Grosh, Lori Diachin, Rob Falgout, and Panayot Vassilevski
GAMS	Solvers	Barry Rountree
GE Energy Consulting	hpc4energy incubator: improving positive sequence load flow simulation performance and capability	Steve Smith
GE Global Research	Wind power	Bill Henshaw
IBM	Advanced memory concepts FastForward project	Bronis de Supinski
IBM	Deep Computing Solutions	Fred Streitz, Jeff Wolf, and Doug East

COMPANY	TOPIC	LLNL CONTACT
IBM	Blue Gene/Q: the evaluation of transactional memory and the modeling	Martin Schulz and Barna Bihari
IBM	Blue Gene/Q common development tools interface co-design	Dong Ahn
IBM	Evaluating the performance of algebraic multigrid on multicore architectures	Ulrike Yang
IBM	Scalable electrophysiological and electromechanical modeling of the human heart	Dave Richards and others
IBM	High performance storage system	Jerry Shoopman
IBM	Scalable systems, multiple areas	Bronis de Supinski
IBM	Tool interface for OpenMP	Martin Schulz
IBM	Flash storage systems	Maya Gokhale and Roger Pearce
IBM and Energy Exemplar	Solvers	Barry Rountree
IBM and Energy Exemplar	Parallel and decomposition schemes for solving large-scale electrical-grid stochastic optimization problems	Deepak Rajan and Carol Meyer
IBM Research	Improvements to CPLEX optimization software geared toward use cases for California Energy Systems for the 21st Century (CES-21)	Deepak Rajan
IBM Research	Operating systems	Maya Gokhale and Scott Lloyd
IBM Research and North Carolina State University	Predictive performance anomaly prevention for virtualized cloud systems	Deepak Rajan
IBM Research, HP Labs, Knight Capital Group, and Bank of America	Scheduling heterogenous jobs in MapReduce environments	Deepak Rajan
IBM Research, Knight Capital Group, StonyBrook University, and Carnegie Mellon University	Mining large time-evolving graphs for proximity queries	Deepak Rajan
InfiniBand Trade Association	InfiniBand specifications body	Pam Hamilton
Intel Corporation	FastForward processor project	Matt Leininger
Intel Corporation	FastForward I/O project	Mark Gary
Intel Corporation	Many integrated core programming environment	Greg Lee
Intel Corporation	Research and development for I/O systems	Mark Gary, Robin Goldstone, and Ned Bass
Intel Corporation	CRADA for system software research and development	Kim Cupps, Pam Hamilton, Chris Morrone, and Martin Schulz
Intel Corporation	Hadoop over Lustre	Al Chu and Robin Goldstone
Intel Corporation and Cray	High performance architecture for data analytics	Matt Leininger, Robin Goldstone, and Trent D'Hooge
ION Geophysical Corporation	Oil exploration	Shawn Larsen
ISO New England and GAMS	hpc4energy incubator: evaluation of robust unit commitment	Deepak Rajan, Barry Rountree, and Liang Min
John Deere	Vehicle simulation and analysis	Jeff Wolf and others
Juelich Research Center/Juelich Supercomputing Center	Tools for performance analysis at scale	Bernd Mohr
Krell Institute/Argo Navis Technologies	Open|SpeedShop development and support and the component-based tool framework	Martin Schulz

COMPANY	TOPIC	LLNL CONTACT
Laboratory for Laser Energetics and Commissariat à l'Énergie Atomique	Miro and virtual beamline modeling and simulation codes	Kathleen McCandless
Life Technologies	Targeted microbial DNA amplification to enhance sequencing	Tom Slezak
Life Technologies	Nucleotide sequence analysis	Tom Slezak and Shea Gardner
Mellanox	Long haul InfiniBand	Trent D'Hooge
Mellanox	Hadoop with InfiniBand remote direct memory access	Al Chu and Robin Goldstone
Micron	Processing in memory	Maya Gokhale and Scott Lloyd
MITRE Corporation	Subsurface modeling	Shawn Larsen
National Instruments	Object-oriented applications of Laboratory view on big physics data	Mike Flegel
NetApp	High performance I/O systems	Marc Stearman and Mark Gary
NSTec	Instrument calibration techniques	Steve Glenn
NVIDIA	FastForward processor project	Bronis de Supinski
OpenFabrics Alliance, Mellanox, and QLogic	OpenFabrics enterprise distribution	Matt Leininger
OpenSFS	Lustre file system development and deployment	Terri Quinn, Pam Hamilton, and Chris Morrone
OpenWorks	Valgrind memory tool and threading tool development	John Gyllenhaal
OpenZFS	ZFS file system development	Brian Behlendorf
OSIsoft	Management and visualization of phasor measurement unit data	Ghaleb Abdulla
Pacific Gas and Electric	Department of Energy Office of Electricity project on supply chain cyber ecurity	Dan Quinlan
ParaTools	Development and support of TAU performance analysis tool	Chris Chambreau
PTC	Windchill	Al Churby
Red Hat	Operating systems	Mark Grondona and Jim Foraker
Red Hat, Appro, Intel, and AMD	Hardware performance counters	Barry Rountree
Roche NimbleGen	Next-generation peptide microarrays	Tom Slezak
Rogue Wave Software	TotalView parallel debugger scalability and enhanced memory tools	Dong Ahn and Scott Futral
Rogue Wave Software	TotalView enhanced debugging for C++ applications	Matt Wolfe
Samplify	Data compression	Peter Lindstrom
San Diego Gas & Electric, Southern California Edison, Pacific Gas and Electric, and the California Public Utilities Commission	California Energy Systems for the 21st Century (CES-21)	John Grosh and Jamie Van Randwyk
ScaleMP	Large memory architectures	Robin Goldstone
SchedMD	SLURM resource management software	Kim Cupps and Don Lipari
STFC Daresbury	National lab collaborations with industry	Doug East, Fred Streitz, and Jeff Wolf
Tennessee Valley Authority and Applied Communication Sciences	Robust adaptive topology control for Advanced Research Projects Agency-Energy project	Barry Smith
Texas A&M Engineering Experiment Station, Arizona State University, UC Berkeley, Applied Communication Sciences, and Tennessee Valley Authority	ARPA-E project: Robust Adaptive Topology Control	Deepak Rajan and Barry Rountree
TidalScale	Bioinformatics applications on virtual large-memory nodes	Alexander Ames
United Technologies Research Center	hpc4energy incubator: improving simulations of advanced internal combustion engines	Steve Smith

NATIONAL LABORATORY COLLABORATIONS

LABORATORY	TOPIC	LLNL CONTACT
Argonne National Laboratory	CESAR Co-Design Center: tools and performance	Martin Schulz
Argonne National Laboratory	Exascale operating systems: power scheduling	Maya Gokhale, Brian Van Essen, Edgar León, Martin Schulz, and Barry Rountree
Argonne National Laboratory	Simulation technologies for multiphysics simulations	Carol Woodward
Argonne National Laboratory and Oak Ridge National Laboratory	CORAL procurement	Bronis de Supinski
Argonne National Laboratory, Brookhaven National Laboratory, Idaho National Laboratory, Lawrence Berkeley National Laboratory, Los Alamos National Laboratory, Pacific Northwest National Laboratory, Sandia National Laboratories	Energy Facility Contractors Group	Darrel Whitney and others
Argonne National Laboratory, Lawrence Berkeley National Laboratory, Los Alamos National Laboratory, Oak Ridge National Laboratory, and Sandia National Laboratories	SDAV: Scalable Data Management, Analysis, and Visualization	Eric Brugger
Argonne National Laboratory, Lawrence Berkeley National Laboratory, Oak Ridge National Laboratory	SUPER: Institute for Sustained Performance, Energy, and Resilience	Marc Casas Guix, Bronis de Supinski, Todd Gamblin, and Daniel Quinlan
Argonne National Laboratory, Lawrence Berkeley National Laboratory, Sandia National Laboratories	FASTMath: Frameworks, Algorithms, and Scalable Technologies for Mathematics	Lori Diachin
Los Alamos National Laboratory	ExMatEx Co-Design Center: tools and performance	Martin Schulz
Los Alamos National Laboratory	Monte Carlo N-Particle Transport Code	Lila Chase
Los Alamos National Laboratory and Sandia National Laboratories	Open|SpeedShop and component-based tool framework	Martin Schulz
Los Alamos National Laboratory and Sandia National Laboratories	Tri-laboratory common computing environment tools	Martin Schulz
Los Alamos National Laboratory and Sandia National Laboratories	TotalView enhanced debugging for C++ applications	Matt Wolfe
Los Alamos National Laboratory and Sandia National Laboratories	Tri-Lab Operating System Software	Pam Hamilton and Trent D'Hooge
Los Alamos National Laboratory, Lawrence Berkeley National Laboratory, Oak Ridge National Laboratory, and Sandia National Laboratories	High Performance Storage System	Jerry Shoopman
Los Alamos National Laboratory, Sandia National Laboratories, and Oak Ridge National Laboratory	ExMatEx Co-design Center: materials in extreme environments	Milo Dorr
Lawrence Berkeley National Laboratory	Extreme resilient discretizations	Jeff Hittinger
Lawrence Berkeley National Laboratory	High-order methods for kinetic simulation of plasmas	Milo Dorr and Jeff Hittinger
Lawrence Berkeley National Laboratory and Sandia National Laboratories	ExaCT Co-Design Center: Center for Exascale Simulation of Combustion in Turbulence	Rob Falgout and Ulrike Yang
Oak Ridge National Laboratory	ExMatEx Co-Design Center: modeling	Martin Schulz
Oak Ridge National Laboratory	Solvers for implicit climate simulations	Carol Woodward and Aaron Lott
Oak Ridge National Laboratory	ARPA-E project: Robust Adaptive Topology Control	Deepak Rajan and Barry Rountree
Oak Ridge National Laboratory and Lawrence Berkeley National Laboratory	SAMRAI performance capabilities	Brian Gunney
Pacific Northwest National Laboratory	Performance analysis for the X-Stack	Martin Schulz
Pacific Northwest National Laboratory	Terrorism risk assessments	Amy Waters and Lisa Belk
Pacific Northwest National Laboratory	Exascale operating systems/runtime	Maya Gokhale, Brian Van Essen, and Edgar León
Sandia National Laboratories	ExMatEx Co-Design Center: Structural Simulation Toolkit	Martin Schulz
Sandia National Laboratories	Multiscale Climate SciDAC: CAM-SE	Aaron Lott
Sandia National Laboratories	Multiscale Climate SciDAC: Trilinos	Aaron Lott
Sandia National Laboratories	Exascale operating systems/runtime	Maya Gokhale and Scott Lloyd

Collaborate
features
found
protect
great
particular
able
threat